

# Java™ magazine

By and for the Java community 

## DevOps

THE PIPELINE TO BUILDING  
BETTER SOFTWARE FASTER

+

28 BUILD A DEVOPS  
PIPELINE WITH  
DOCKER AND PUPPET

22 HOW NETFLIX  
RUNS ITS DEVOPS  
PIPELINES

40 FUNCTIONAL  
PROGRAMMING IN  
JAVA 8

## 01

## EDITORIAL

### Editor in Chief

Andrew Binstock

### Senior Editor

Claire Breen

### Copy Editor

Karen Perkins

### Section Development

Michelle Kovac

Bob Larsen – Java Nation

Yolande Poirier – Java Nation

### Technical Reviewers

Stephen Chin, Reza Rahman, Simon Ritter

## DESIGN

### Senior Creative Director

Francisco G Delgadillo

### Design Director

Richard Merchán

### Contributing Designers

Jaime Ferrand, Arianna Pucherelli

### Senior Production Manager

Sheila Brennan

### Production Designer

Kathy Cygnarowicz

## PUBLISHING

### Publisher

Jennifer Hamilton +1.650.506.3794

### Associate Publisher and Audience Development Director

Karin Kinnear +1.650.506.1985

### Audience Development Manager

Jennifer Kurtz

## ADVERTISING SALES

### President, Sprocket Media

Kyle Walkenhorst +1.323.340.8585

### Western and Central US, LAD, and Canada, Sprocket Media

Tom Cometa +1.510.339.2403

### Eastern US and EMEA/APAC, Sprocket Media

Mark Makinney +1.805.709.4745

### Advertising Sales Assistant

Cindy Elhaj +1.626.396.9400 x 201

### Mailing-List Rentals

Contact your sales representative.

## RESOURCES

### Oracle Products

+1.800.367.8674 (US/Canada)

### Oracle Services

+1.888.283.0591 (US)

### Oracle Press Books

[oraclepressbooks.com](http://oraclepressbooks.com)

## ARTICLE SUBMISSION

If you are interested in submitting an article, please [e-mail the editors](#).

## SUBSCRIPTION INFORMATION

Subscriptions are complimentary for qualified individuals who complete the [subscription form](#).

## MAGAZINE CUSTOMER SERVICE

[java@halldata.com](mailto:java@halldata.com) **Phone** +1.847.763.9635

## PRIVACY

Oracle Publishing allows sharing of its mailing list with selected third parties. If you prefer that your mailing address or e-mail address not be included in this program, contact [Customer Service](#).

**Copyright © 2015, Oracle and/or its affiliates.** All Rights Reserved. No part of this publication may be reprinted or otherwise reproduced without permission from the editors. *JAVA MAGAZINE* IS PROVIDED ON AN "AS IS" BASIS. ORACLE EXPRESSLY DISCLAIMS ALL WARRANTIES, WHETHER EXPRESS OR IMPLIED. IN NO EVENT SHALL ORACLE BE LIABLE FOR ANY DAMAGES OF ANY KIND ARISING FROM YOUR USE OF OR RELIANCE ON ANY INFORMATION PROVIDED HEREIN. The information is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle. Oracle and Java are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

*Java Magazine* is published bimonthly at no cost to qualified subscribers by Oracle, 500 Oracle Parkway, MS OPL-3A, Redwood City, CA 94065-1600.

Digital Publishing by [GTxcel](#)



# 20 Years of Innovation

## #1 Development Platform



Since 2002



Since 1999



Since 1995



Since 1996



Since 1998



Since 1996



Since 1996



Since 2008



Since 1998



Since 1999



Since 2001



Since 1996



Since 2012

# ORACLE®



A man with glasses, wearing a light blue button-down shirt and blue jeans, is walking towards the camera on a city street. He is holding a grey folder or tablet under his left arm. The background is a blurred city street with other pedestrians and buildings.

# Primitive Obsession: A Little-Discussed but Expensive Antipattern

**M**ost code smells, it seems, develop incrementally over time. Code that starts out as good goes bad by a series of subsequent maintenance passes that are not coupled with intensive refactoring. For example, smells such as long parameter lists, long functions, and God objects frequently are the product of slowly accreted functionality. Except at organizations that rigorously do code reviews, these problematic bits can live a long time in the codebase, bringing down its quality and building up technical debt.

One antipattern that survives longer than most is *primitive obsession*. This term refers to different things depending on whom you ask, but all definitions share one aspect: the use of a raw data type (a

A classic example is using an integer to hold a US zip code (generally, a five-digit number) in a method that, say, prints addresses. If later, the program begins accepting the four-digit zip code extensions, new formatting logic will be required, because the printed form is no longer just a string of numbers. If the program then adds a zip-based barcode, additional logic will be needed. If the source code, which has been revised multiple times, remains in the original address-printing method, things get cumbersome. Imagine now, if you wanted to add postal codes for the UK, which include letters. The primitive obsession would become so obvious that the natural refactoring would be undertaken: Replace

CREATE  
THE FUTURE

oracle.com/java

20  
YEARS  
1995-2015

Java™

ORACLE®



```
//from the editor /
```

data value with object.

As a fan of small classes, I am sensitive to the opportunity to get rid of primitive obsession. Anytime a local field accretes logic specific to its operation or formatting, I move it into its own class. This step makes the parent class smaller and less complex; it also enforces the Single Responsibility Principle, which is one of the keys to effective object orientation. The new class is also easier to test and verify. And because the class now hides implementation details, changing the internal representation from an integer to a string, when the UK addresses are added, is a modification of comparatively small scope.

This data hiding is a signal benefit when the data item is a collection rather than a primitive. It can hide access to the collection's elements. For example, if I were to use a naked Java collection, rather than a wrapper object, any other method could come along and insert elements or even run `reverse()` or `swap()` without my say-so. With my wrapper, I can specify exactly what can be done to the data and so prevent all kinds of mischief.

When I return a data item, it's almost invariably immutable; and in the event an element that doesn't exist is asked for, I return an empty string, or an empty array, rather than a null. These steps mean I'm pumping cleaner data and fewer potential NPEs into the codebase. This is a lot of benefit for a simple operation.


So, if you have a data item (primitive or collection) that begins to accrete logic, move it to its own object. You'll be glad you did.

**Andrew Binstock, Editor in Chief**  
[javamag\\_us@oracle.com](mailto:javamag_us@oracle.com)  
[@platypusguy](#)

PS: Our continuing migration to a more technical, more engaging, and less promotional magazine continues. Code listings are now inline in the articles, which should greatly increase readability; we've also added a section on letters to the editor; and the books section now consists entirely of reviews, rather than promotional descriptions. Our next issue will be larger, add more innovations, and show the vision of what we want *Java Magazine* to be.

# Prove Your Tech Creds

# Get Java Certified



## Save up to 20%

- ✓ **Get noticed by hiring managers**
- ✓ **Learn from Java experts**
- ✓ **Join online or in the classroom**
- ✓ **96% of participants recommend it**
- ✓ **100% report promotions, raises, and more**



# JRebel

---

 ZEROTURNAROUND

RELOAD CODE CHANGES  
INSTANTLY

**TRY IT NOW!**

*Get a free  
t-shirt! →*





MAY/JUNE 2015

## Listings Are Hard to Read

Since you asked for feedback: the listings in the articles are a nightmare. The pages (even in the PDF download) show only the Listing 1, and when I click on the other listings I'm redirected to some odd web page that either doesn't open or redirects to some general download page where I need to download all listings and search for the one I'm interested in. Takes me a minute, and by then I have forgotten what I read. The listings have to be beside the text—any other solution is complicated.

Things get worse if I'm offline: the listings in the PDF tabs don't work.

—Christian Schenk

*Editor Andrew Binstock responds: Your frustration was echoed in other readers' comments and in my own experience as a longtime reader of Java Magazine. In this issue, we've begun to address the problem of listings. To the extent possible, code is now presented inline in the articles. There will be no tabs to click on.*

One exception will be listings that are very long. These you'll need to download. With that in mind,

*we've renovated the download area. Listings are no longer placed in one file, but rather presented as individual code listings in plain-text files.*

*If you have any further suggestions or comments on these changes, please let me know. Feedback is always welcome. Information on how to contact me is available on the last page of this issue.*

## Confusion Over the Java char Data Type

Dear Editor,

Your review of the book *Core Java for the Impatient* by Cay S. Horstmann on page 11 of the May/June 2015 issue of *Java Magazine* has misinterpreted a Java fact. You noted a contradiction between the following two statements in the book: "Java has ... five integral types" and that in Java "four [primitive types] are integer types."

These are actually not contradictory statements. What the review has not taken into consideration is the distinction between “integral type” and “integer.” While `char` is an integral type, its value is a Unicode character and not an integer. Refer to the following two ex-

cerpts from the Java 8 specification.

“The primitive types are defined to be the same on all machines and in all implementations, and are various sizes of two’s-complement integers, single- and double-precision IEEE 754 standard floating-point numbers, a boolean type, and a Unicode character `char` type.”

“The integral types are **byte**, **short**, **int**, and **long**, whose values are 8-bit, 16-bit, 32-bit, and 64-bit signed two’s-complement integers, respectively, and **char**, whose values are 16-bit unsigned integers representing UTF-16 code units.”

Notice the distinction between integral types `byte`, `short`, `int`, and `long` and the integral type `char`. Only four integral types (`byte`, `short`, `int`, and `long`) are two's-complement integers. The `char` type is an integral type but not a two's-complement integer; instead, `char` is a Unicode character.


I would suggest that an erratum be included in the next issue of *Java Magazine* about the book review being a misinterpretation of a Java fact.

—Deepak Vohra









The well-written, approachable text is supported with sensible

To use this book, you will need WildFly 8.1 (Red Hat's renamed

This book can be recommended only to readers who are searching for very specific WildFly-based recipes, who can work strictly in Eclipse, and who are willing to overlook its errors. —A.B.



# DevOps: It's Not Just For The Unicorns

There's a stigma that exists where we constantly see suggestions that DevOps is just for the unicorns. This couldn't be more false. Afterall, those unicorns were once horses themselves.



Ready to start your journey to becoming a successful DevOps IT organization?

Download your copy of "Making The Shift: From Continuous Integration to Continuous Delivery" and start automating your development and delivery pipeline.



**CloudBees**  
The Enterprise Jenkins Company





Santa Clara, California

## **JVM Language Summit** AUGUST 10–12

SANTA CLARA, CALIFORNIA

The JVM Language Summit is an open technical collaboration among language designers, compiler writers, tool builders, runtime engineers, and VM architects.

## **JCertif 2015**

AUGUST 14–16; KINSHASA,

REPUBLIC OF THE CONGO

SEPTEMBER 7–13;

BRAZZAVILLE, CONGO

JCertif International has trained more than 5,000 new Java developers in Africa since it began, and its conferences are expanding to more countries. The events usually include two- to five-day Java workshops for 20 to 50 develop-

ers. This year's topics include Java; Android; and wearable, object-connected, and drone technologies.

## **JavaZone**

SEPTEMBER 9–10

OSLO, NORWAY

JavaZone, created by the Norwegian Java User Group (javaBin), consists of almost 200 speakers, 2,500 attendees, and seven parallel sessions. An additional day of workshops beforehand is included in the cost of the ticket.

## **JDayLviv**

SEPTEMBER 19

LVIV, UKRAINE

Enjoy and learn in a full day of world-class talks. Topics include JVM languages, cloud and infrastructures, big data and high load, architecture, and front end.

## **code.talks 2015**

SEPTEMBER 29–30

HAMBURG, GERMANY

More than 1,500 participants take part in 110 presentations across 13 tracks, including Java, PHP, JavaScript, DevOps, big data, e-commerce, mobile, and more.

## **Silicon Valley Code Camp**

OCTOBER 3–4

SILICON VALLEY AREA,  
CALIFORNIA

Last year, 4,500 people attended this free community event where developers learn from fellow developers. In addition to technical topics, speakers present on software branding and legal issues.

## **JAX London**

OCTOBER 12–14

LONDON, ENGLAND

JAX London brings Java, JVM, and enterprise professionals together for a technology- and methodology-packed event. Participants get full access to Big Data Con London, which features modern datastores, big data architectures based on Hadoop, and advanced data processing techniques.

## **JDD**

OCTOBER 12–13

KRAKOW, POLAND

JDD is a two-day conference for all Java enthusiasts, who can participate in more than 30 lectures, workshops, interactive trainings, and networking opportunities. JDD attracts speakers from all over





the world. It also offers lectures in English.

### GeeCON

OCTOBER 23–24

PRAGUE, CZECH REPUBLIC

Join more than 2,000 participants at GeeCON, which is focused on JVM-based technologies, with special attention to dynamic languages such as Groovy and Ruby. GeeCON participants share experiences about development meth-

odologies and craftsmanship, enterprise architectures, design patterns, distributed computing, and more.

### JavaOne 2015

OCTOBER 25–29

SAN FRANCISCO, CALIFORNIA

Join the single largest gathering of Java developers. From sessions, workshops, labs, and exhibits to keynotes and Birds-of-a-Feather sessions, learn about the latest

language changes to improve coding efficiency. You'll also learn how to build modern enterprise and server-based applications, create rich and immersive client-side solutions, build next-generation apps targeting smart devices, and compose sophisticated Java web services and cloud solutions.

### W-JAX 15

NOVEMBER 2–6

MUNICH, GERMANY

The W-JAX conference covers current and future-oriented technologies from Java, Scala, Android, and web technologies to agile development models and DevOps.

### Devoxx Belgium

NOVEMBER 9–13

ANTWERP, BELGIUM

By developers for developers, this event has 200 speakers and 3,500 attendees from 40 countries. Tracks this year include Java SE, JVM languages, and server-side Java, as well as cloud and big data, mobile, and architecture and security, among others.

### Devoxx Morocco

NOVEMBER 16–18

CASABLANCA, MOROCCO

Formerly the JMaghreb conference, this event is a university day

of training, workshops, and labs followed by conference days of sessions on software development, web, mobile, gaming, security, methodology, Internet of Things (IoT), and cloud. The Decision Makers evening includes discussion of issues related to the IT industry in Morocco.

### Groovy and Grails eXchange 2015

DECEMBER 14–15

LONDON, ENGLAND

Stay ahead of the curve and hear the 2016 roadmap for Groovy and Grails from core committers and Groovy authorities Guillaume Laforge and Graeme Rocher. Engage with other leading experts and fellow enthusiasts and learn the latest innovations and practices.

Have an upcoming conference you'd like to add to our listing? Send us a link and a description of your event at least four months in advance at [javamag\\_us@oracle.com](mailto:javamag_us@oracle.com). We'll include as many as space permits.



# Reach More than 800,000 Oracle Customers with Oracle Publishing Group

**Connect with the Audience that  
Matters Most to Your Business**



## **Oracle Magazine**

**The Largest IT Publication in the World**

Circulation: 500,000

Audience: IT Managers, DBAs, Programmers and Developers



## **Profit**

**Business Insight for Enterprise-Class Business Leaders to  
Help Them Build a Better Business Using Oracle Technology**

Circulation: 100,000

Audience: Top Executives and Line of Business Managers



## **Java Magazine**

**The Essential Source on Java Technology, the Java Programming  
Language and Java-Based Applications**

Circulation: 200,000 and Growing Steady

Audience: Corporate and Independent Java Developers,  
Programmers, and Architects



# Quiz Yourself

*You never create technical debt. It's just that other developers don't really know the details of how Java works. We completely understand.*

The questions in this quiz section are taken from the certification test 1Z0-808: Oracle Certified Associate, Java SE 8 Programmer, Oracle Certified Java Programmer. (The purpose of this certification is to enable beginners to demonstrate their knowledge of fundamental Java concepts. The certification is designed for beginners in programming and/or those with a nonprogramming background who have basic mathematical, logical, and analytical problem-solving skills and who want to begin to learn the Java programming language; and for novice programmers and those programmers who prefer to start learning the Java programming language at an introductory level. After successfully completing this certification, candidates can confidently utilize their knowledge on Java datatypes, operators, statements, arrays, lists, and exception-handling techniques.)

Naturally, these questions have small embedded traps that spring open if you rush through without considering exactly what the code is doing. Ready? (Answers appear in the “Solutions” section after the questions.)

**Question 1. Given this code fragment:**

```
List<String> vitamins = new ArrayList<>();
vitamins.add("A");
vitamins.add("B12");
vitamins.add("C");
vitamins.set(1,"B");
vitamins.add(1,"D");
System.out.println(vitamins);
```

## What is the result?

- a.**  $[A, D, B, C]$

- b.** [A, D, C]  
**c.** [D, B, C]  
**d.** [D, B12, C]

Question 2. Given this code fragment:

```
1. class Engine { }
2. public class App {
3.     public static void main(String[] args) {
4.         Engine e = new Engine();
5.         Engine e1 = e;
6.         e = null;
7.     }
8. }
```

Which statement is true about this code?

- a.** It creates an object and the object is eligible for garbage collection.
- b.** It creates an object and the object is *not* eligible for garbage collection.
- c.** It creates two objects: **e** and **e1**. The **e** object is eligible for garbage collection.
- d.** It creates two objects and both the objects are *not* eligible for garbage collection.

Question 3. Given this code fragment from `Cart.java`:

```
package shop;
// line n1
{
    Cart() {
        System.out.println("Cart is created.");
    }
}
```



```
//fix this /
```

}  
}

And this content from Shop.java:

```
package shop;
public class OnlineCart extends Cart{
    public static void main(String[] args) {
        Cart c = new OnlineCart();
    }
}
```

Which code fragment can be inserted at **line n1** to enable the **Shop** class to compile?


- final class Cart
- public class Cart
- private class Cart
- protected class Cart

Question 4. Given this code fragment:

```
double wage = 2.0;           // line n1
int weekDays = 5;
long monDays = weekDays * 4; // line n2
long yearDays = monDays * 12L; // line n3
long totalWage = yearDays * wage; // line n4
```

**Which modification enables the code to compile?**

- Replace line n1 with `double wage = 2;`
- Replace line n2 with `long monDays = weekDays * 4L;`
- Replace line n3 with `long yearDays = monDays * 12;`
- Replace line n4 with `long totalWage = yearDays * (long) wage;`



**Question 1.** The correct answer is A. Explanation:

`java.util.List` is an index-based collection. The `set()` method replaces the element at the specified position in the list with the specified element. The `add()` method inserts the specified element at the specified position in the list.

**Question 2.** The correct answer is B. Explanation:

At line 4, the code creates an object that is accessible by using the `e` reference variable.

At line 5, **e1** is another reference variable for the same object and there are two reference variables for the **Engine** object. Hence, the object is accessible by using both the **e** and **e1** variables.

At line 6, the assignment of `null` to the `e` reference variable makes the object inaccessible through `e`.

Option A is incorrect because the object is accessible through the **e1** reference variable. The inaccessible object is eligible for garbage collection. Options C and D are incorrect because only one object is created.

**Question 3.** The correct answer is B. Explanation:

Option A is incorrect because a **final** class cannot be inherited. Option B is correct because the **public** classes are accessible and inheritable. Options C and D are incorrect because the top-level class, **Cart**, cannot be declared with a **private** or **protected** access modifier.

**Question 4.** The correct answer is D. Explanation:

Option D is correct. Line n4 fails to compile due to loss of precision. To fix the error, the `wage` variable must be cast to the `long` type. Option A is incorrect. It is a valid assignment, hence the modification suggested in Option A does not meet the requirement. Options B and C are incorrect. They are valid assignments. The suffix `L` converts to `long` type.

*[Our gratitude to Sushma Jagannath of Oracle University for her kind help in providing these test extracts. —Ed.]*




# TeamCity builds your day

Hassle-free continuous integration  
and deployment server by JetBrains

*Try free. Use free.*

**START BUILDING TODAY**





# DevOps: BUILDING THE PIPELINES

When conceived by its creator, Patrick Debois, DevOps was intended to bring developers and operations (that is, system administrators) into closer cooperation. As he states in our interview ([page 18](#)), DevOps has given developers more insight into how their code runs and helped admins become less risk-averse. As Netflix shows ([page 22](#)), substantial investments in a comprehensive DevOps pipeline can provide those benefits and other unexpected insights. However, when we look at details of pipeline management by using Docker ([page 28](#)) or Chef ([page 35](#)), it's

clear that for small to midsize organizations, DevOps can represent a lot of new work, particularly for developers. A cloud orientation will help, as the needed pipelines are easier to set up and manage for devs and admins.

This cloud dimension pervades not only these articles but DevOps at large. The transition from continuous delivery, as a bleeding-edge extension to continuous integration, to DevOps maps almost directly to the transition from one-off virtualization to the large-scale cloud as a platform. In this sense, DevOps is completely new. —*Andrew Binstock*





Patrick Debois (middle), who coined the term *DevOps*, currently works in a consultancy in Ghent, Belgium.



DevOps

# HOW DEV VS. OPS BECAME DEVOPS

The founder of DevOps explains how the movement started and where it's headed.

BY STEPHEN CHIN

**W**idely known as the father of the DevOps movement for his work in bootstrapping the influential *DevOpsDays* conferences, Patrick Debois recounts the origins of the term, the concepts, and the subsequent growth of what is now standard practice at many large developer IT organizations.

**Java Magazine:** Tell us a little bit about how you came to found DevOps.

**Debois:** Six or seven years ago, I was working on a project that involved agile development. Although I was a database admin, I really liked the flexibility of the agile mentality and tried to translate everything that agile development was doing into my systems work—such as faster feedback, testing at the infrastructure level, and so forth. We called what we were doing at the time “agile system administration.”

PHOTOGRAPHS BY TON HENDRIKS





**Java Magazine:** Not nearly as cool sounding as “DevOps.”

**Debois:** It was also narrowly focused, because it only included database and system administration. Then I went to the 2008 Agile Toronto conference and met Andrew Shafer [[@littleidea](#) on Twitter], a software developer who was already talking about agile infrastructure. We started the Agile System Administrator Google Group. Later, I saw a talk by [Jean-Paul Sergent](#) about Dev and Ops working together. I didn't really think about the term DevOps at that time, but I mentioned on Twitter that it would be cool to have a conference in Europe exploring some of these ideas. "Agile System Administrators" isn't a very good name for a conference, so Sergent suggested we call it "DevOpsDays." And that is how the DevOps movement started.

**Java Magazine:** So the conference name came before the term *DevOps*?

**Debois:** Yes. We used Twitter to connect with about 60 people in Belgium. Two things were memorable about that conference. One was the energy of the conference people. Secondly, it was very international. Several attendees flew in from Australia and the US. Others came from France, Germany, and the UK. We started with just a couple of tweets about something that didn't exist, and we had 60 people from all over the world come to that conference. After the conference, we continued our discussions using the hashtag #DevOps.

**Java Magazine:** Despite your initial Ops perspective, a lot of DevOps traction

“If you’re not willing to collaborate,  
it doesn’t really matter what tool  
you use or don’t use.”

has come from developer interest. What attracts developers to DevOps?

**Debois:** I came from the admin side, but I always had an integration role. In the early days when Java didn't have servlets, I created my own servlet loaders because I needed to have functioning websites. I had an affinity for developers because it

always struck me that when you're in the developer role, you often hate the operational constraints. In discussions about infrastructure, admins would throw comments at you, such as, "No, it's not the correct directory. You need to change the ports." It's really annoying that you can't handle that yourself.

**Java Magazine:** Hence the need for Dev and Ops integration.

**Debois:** Before the DevOps movement, once your app was running, you didn't have any insight into what else was going on. You couldn't see the logins, you didn't have access to the file system if something failed, and you couldn't do debugging. You couldn't see how things ran in production in order to avoid making the same mistakes. Those were the main drivers that led me to believe that all this needs to be more flexible. Why can't developers and operations collaborate instead of hitting a brick wall? That was the incentive most developers I worked with had, and I think most people have.

If you're really passionate about what you're building and you want to improve things, you do want to go that extra mile, and you do want to have that information, and you do feel responsible about how your application is doing in production.

**Java Magazine:** Much of your focus has been on how people interact, as well as streamlining how developers and operations work together. In the past five years, there have been a lot of advances in tooling associated with DevOps. Are these tools changing the way that people do development, or have they become too much of a crutch?

**Debois:** There is always pushback when you try to change culture, and a tool by itself can't change culture. What's good



“In general, a developer assumes the world doesn’t change outside his app, and an admin assumes the opposite.”

A good example of this is virtualization. In the beginning, people were using it in much the same way as when they didn't have virtualization. But once you start thinking deeply about virtualization as disposable and rebuildable environments, your concept of virtualization changes. One of the things that good tools can do is give you faster feedback on things that you or someone else is doing—feedback that can be captured via a simple chat system or as part of your testing system.

Tools that improve feedback also boost collaboration. You might want to collaborate, but if it takes a week to get a report about a problem, you feel less incentive to work on it. It's what's in your immediate attention span that keeps you going. The risk, of course, is if you're doing continuous integration or continuous delivery, you take it to a point where you're not getting feedback from your customers.

**Java Magazine:** That's pushing it to the extreme if you're not getting feedback from your customers.

**Debois:** That is an agile concept that was missing in practice. Operations was still finding anomalies and saying somebody else in the pipeline was responsible, which shouldn't be the case. Tools that get information to the developers, such as dashboards and monitoring systems, all help provide feedback and shorten the cycle. They support your process tremendously if you are willing to collaborate. Of course, if you're not willing to collaborate, it doesn't really matter what tool you use or don't use.

**Java Magazine:** How has DevOps changed the views of admin teams?

**Debois:** DevOps has changed the culture in the admin world. Admins used to be risk-averse. Their mantra used to be “we don’t do this.” Now, admins are saying “let’s do it more often if it’s hard,” “let’s see what happens when we try to automate,” or “let’s feel the pain and then discuss things because

now we're an equal stakeholder in the discussion." For developers, who have always been at the center of feature discussions and product decisions, that culture was already there if they were doing agile. But within an operations group or a systems group, this was not a common mentality.

**Java Magazine:** In the early days of extreme programming, we were pushing the development envelope and people were averse to making changes because they worried that more-frequent releases would introduce more bugs. Now, it's completely shifted in the other direction to where most organizations are embracing agile development. It sounds as though DevOps is having the same kind of cultural effect on operations groups by making drastic changes in how applications are deployed.

**Debois:** This change comes partly from the fact that, in modern infrastructures, we're dealing less with hardware and more with an API or programmable infrastructure. So admins have evolved to thinking more programmatically about infrastructure, which is a natural fit with the current development process. Feedback cycles are often still too slow: testing on Ubuntu and creating a VM [virtual machine] is much slower than doing something locally, although technologies such as Docker are improving the feedback cycle.

But it's complex. In general, a developer assumes the world doesn't change outside his app, and an admin assumes the opposite: he owns the app and has to keep it running. An admin doesn't have the skills to change the app, but he knows the world is changing and so he deals with that part of the world that involves protecting the app or making sure something is patched.

**Java Magazine:** Are there other disciplines where you can apply the best practices and experience you've gained from DevOps?

**Debois:** Interesting question. I deliberately made the decision of not working for another big technology company. I didn't want to be in the back room; I wanted to be in on product discussions. I realized that I could use all the technology knowl-





# CONTINUOUS DELIVERY OF MICROSERVICES AT NETFLIX

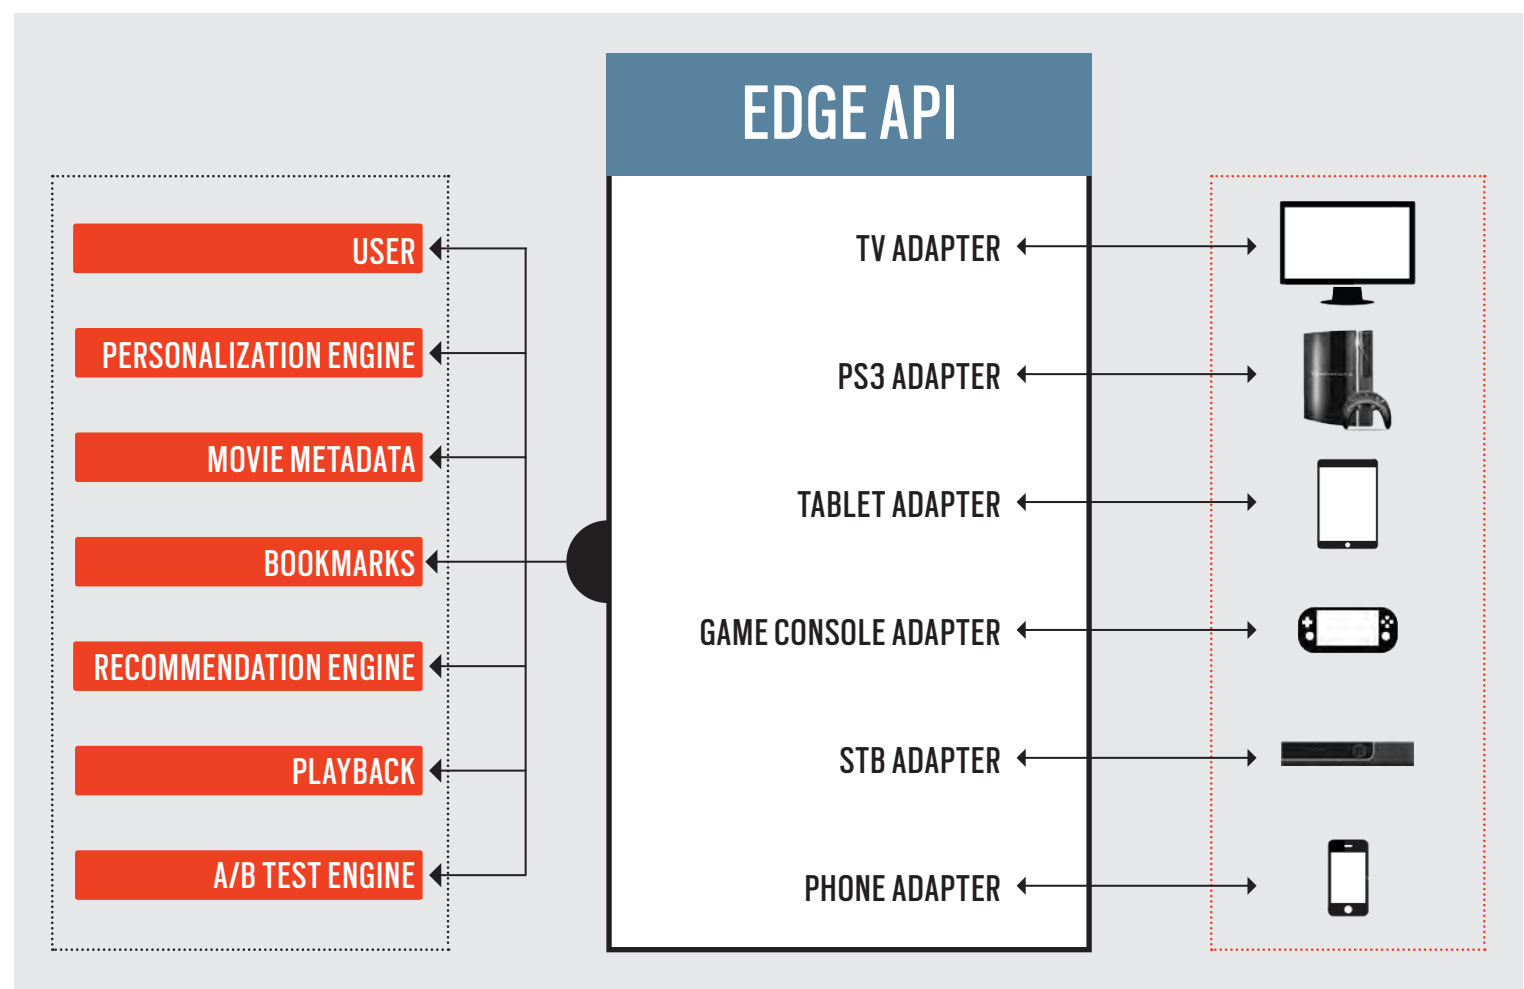
By innovating its development pipeline to accommodate microservices, Netflix automates updates and rapidly tests new ideas. **BY SANGEETA NARAYANAN**

In today's marketplace, companies need to innovate continuously in order to stay relevant and sustain growth. Consequently, velocity becomes a primary requirement for software engineering organizations. To move fast, systems need to be architected with agility in mind. For such systems, the actual process of software delivery can be fully automated, as we have done at Netflix by using continuous delivery to roll out new versions of our microservices.

*Microservices* are an architectural approach in which a single application is built from a suite of small, collaborating services. Each service is responsible for a subset of the application functionality and can be operated independently of the others.

*Continuous delivery* (CD) is a software development practice that makes it cheap, quick, and easy to roll out new versions of an application with confidence. The idea is to develop a software delivery process that allows every commit to automatically be deployed to production, with visibility into the entire process.

A combination of microservices and CD enables companies such as Netflix to rapidly test new ideas and continuously improve the customer experience.



**Figure 1. The Netflix Edge API Service**





## THE NETFLIX EXAMPLE

The Netflix product consists of a set of microservices running on the Amazon Web Services (AWS) cloud, each focused on a specific business function. Subscribers worldwide watch Netflix on more than 1,000 device types. The Netflix Edge API service powers the user experience by connecting subscribers' devices to the Netflix streaming ecosystem in AWS.

The API, like a majority of Netflix microservices, is a Java application executing within an Apache Tomcat container. Communication between devices and the API occurs over HTTP, as does the interaction between the API and the other services.

**Figure 1** provides an overview of the Netflix streaming service. Each box in this image represents a component that is developed and operated independently. This decoupling is a key enabler of agility. The sections that follow touch upon a few CD concepts using the API process as an example.

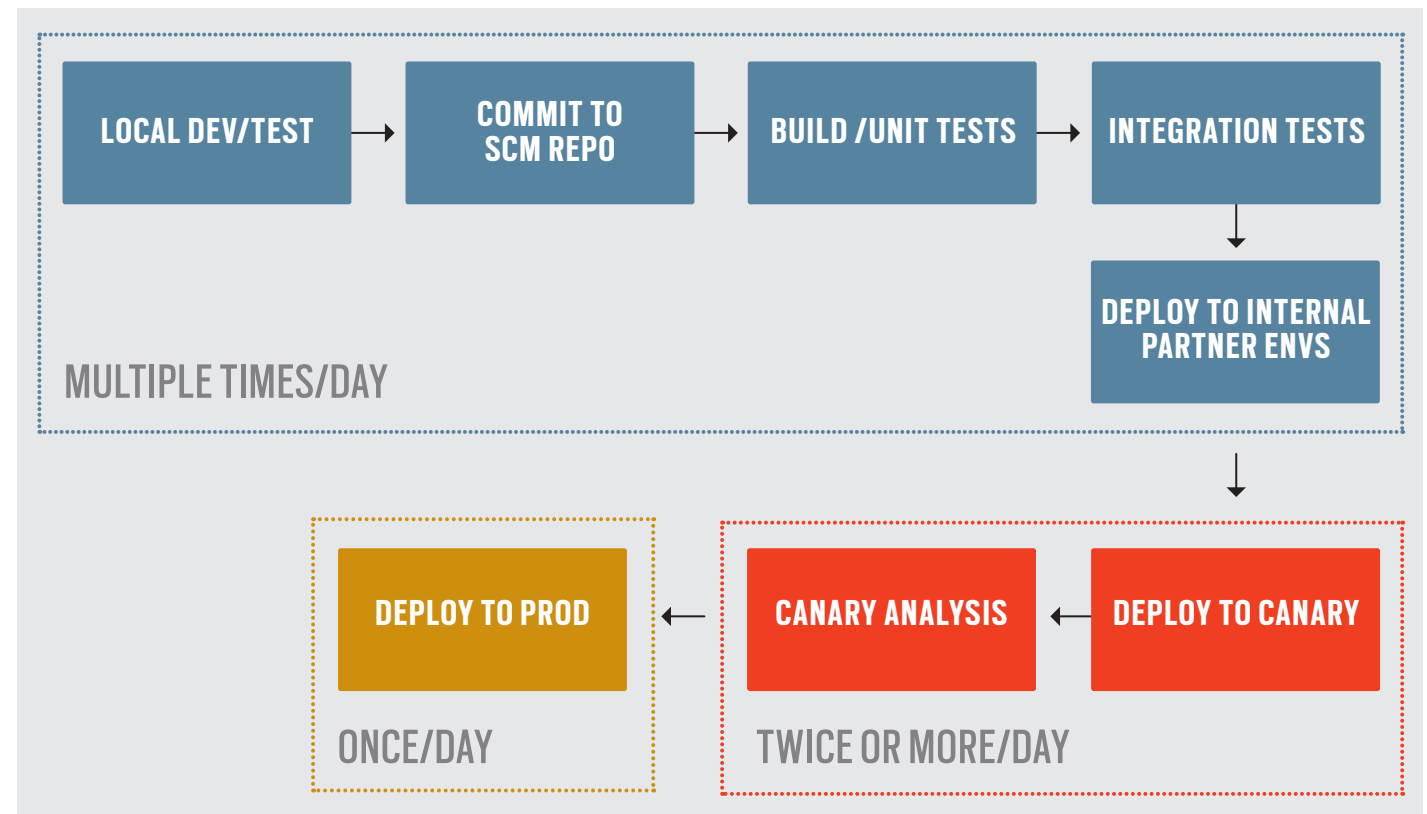
## THE NETFLIX EDGE API DELIVERY PIPELINE

Delivery pipelines are an important concept in the practice of CD. A pipeline consists of stages through which code passes, eventually resulting in a deployable artifact. Each stage provides feedback about code quality and deployment readiness. A failure in any stage halts the pipeline, and no code can move through the system until the failure is corrected.

The API delivery pipeline is illustrated in **Figure 2**.

**Development and testing.** Every commit starts out as a release candidate, so it is important to select a development workflow optimized for CD. The API workflow is modeled after [GitHub Flow](#), with one variation: pull-request processing has been automated to run a test suite in addition to allowing for optional manual code reviews. This ensures that the code won't break the master build.

Automated tests are the foundation of a delivery pipeline. For us, unit tests and integration tests—coupled with automated canary analysis (discussed later)—form the core of the test strategy for the Netflix Edge API.



**Figure 2.** API Delivery Pipeline

**Dependency management.** Dependency management is an important consideration when building Java applications. The ideal dependency management strategy balances the need for generating reproducible builds while allowing for rapid dependency updates.

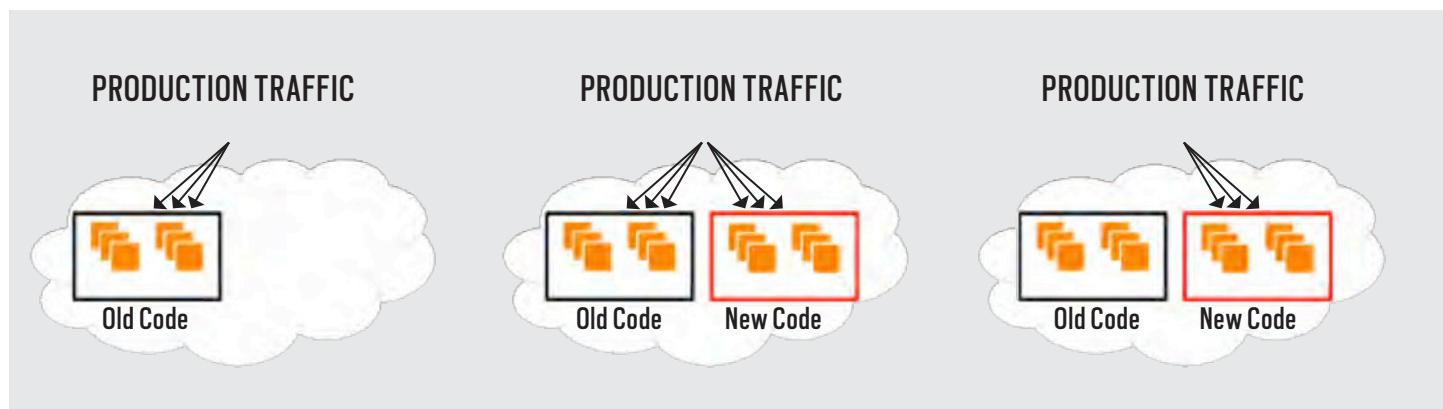
The Netflix Edge API has a complex dependency graph based on hundreds of libraries, most of which are being updated continuously. [Gradle](#) is used for build automation and dependency management, along with a few custom plugins that have been open sourced under [Nebula](#). The [gradle-dependency-lock](#) plugin resolves dependencies based on dynamic versions and allows locking of dependency versions based on the resulting graph. Here is a snippet from a [build.gradle](#) file that is configured to always pull the latest released version of its dependencies:











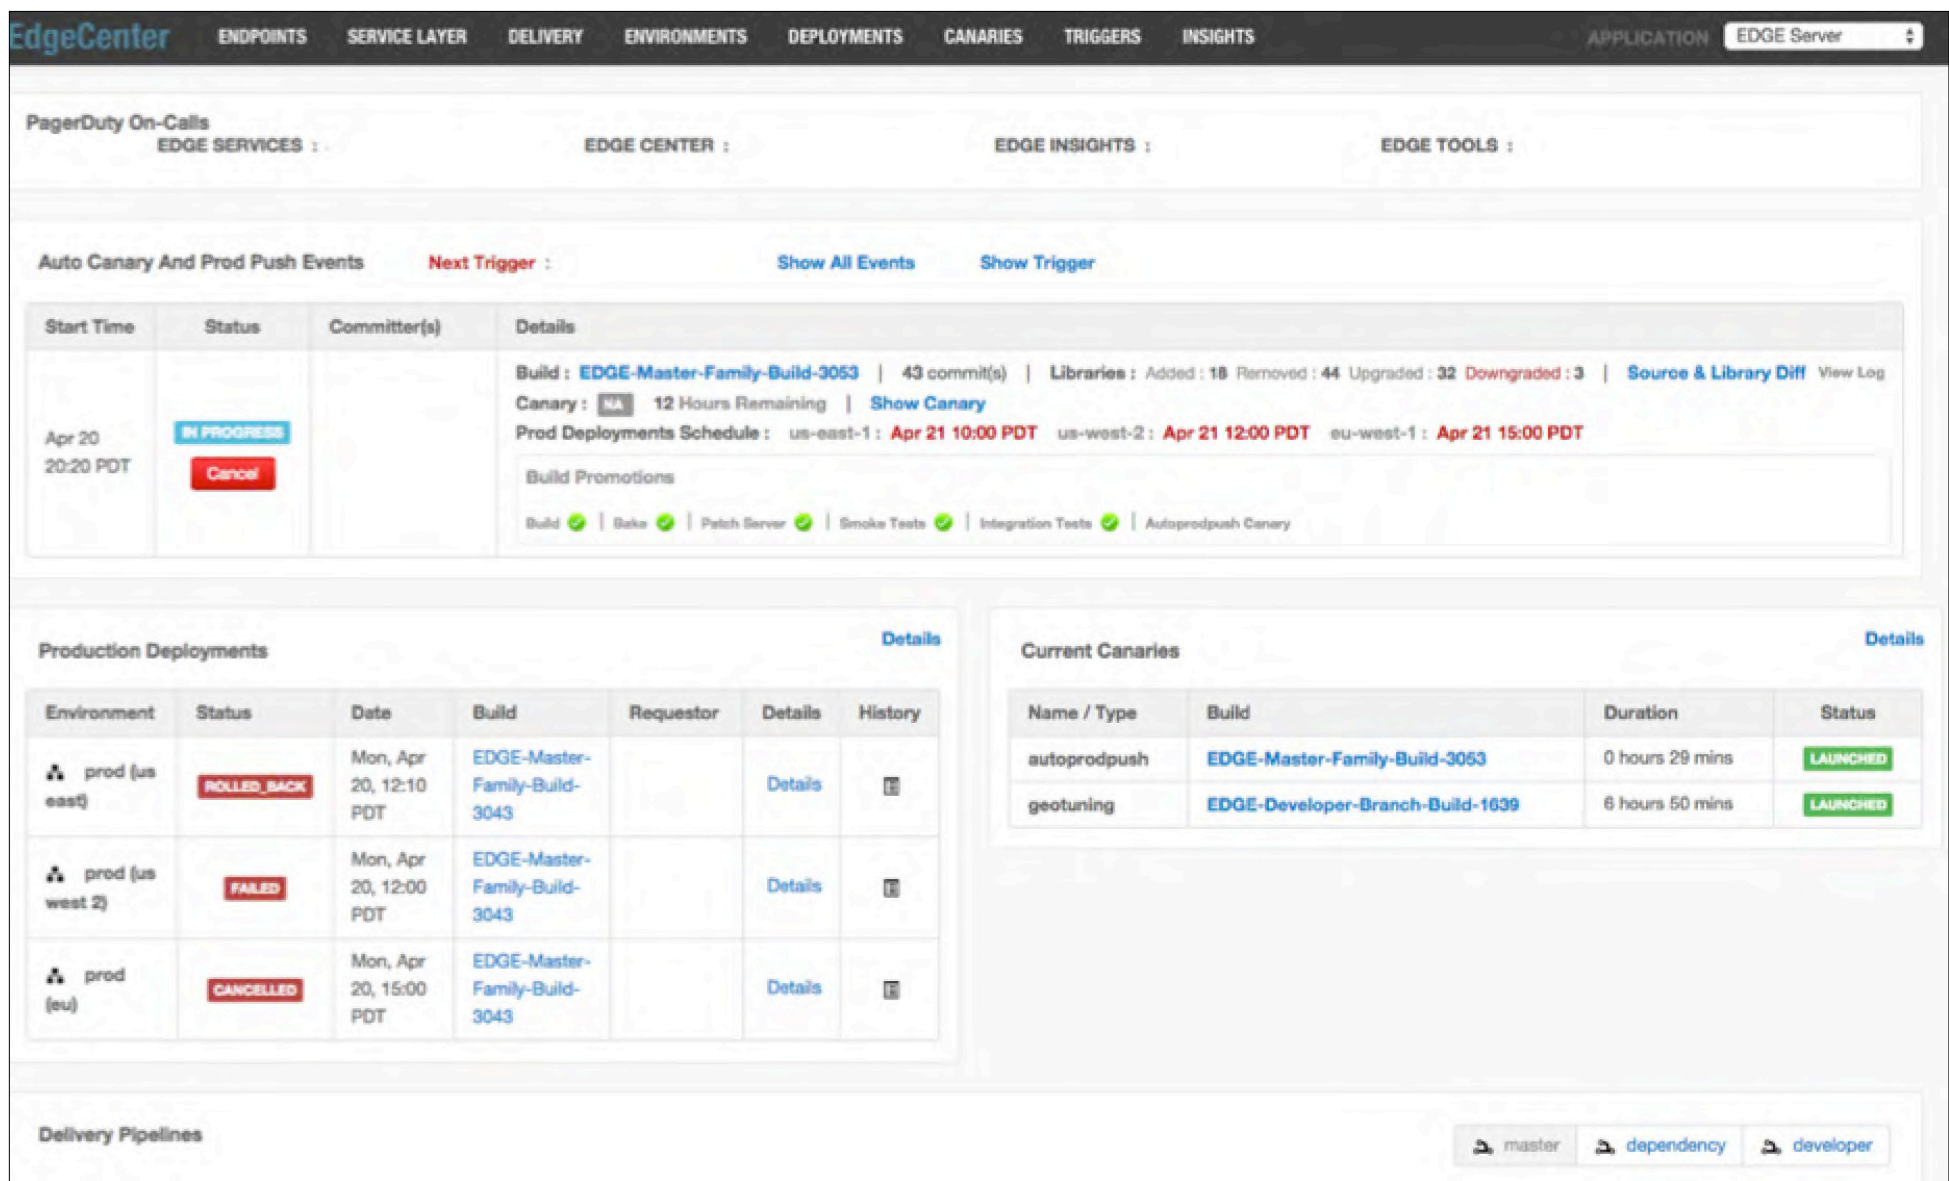
tion are published to an internal event tracking system for auditability.

Microservices and CD pave the way for a distributed model for operations. Netflix engineering follows the “Operate What You Build” model. API developers are responsible for the development, testing, and operations of their code. They engage with business users and manage the rollout of the features they own. Additionally, they manage canary

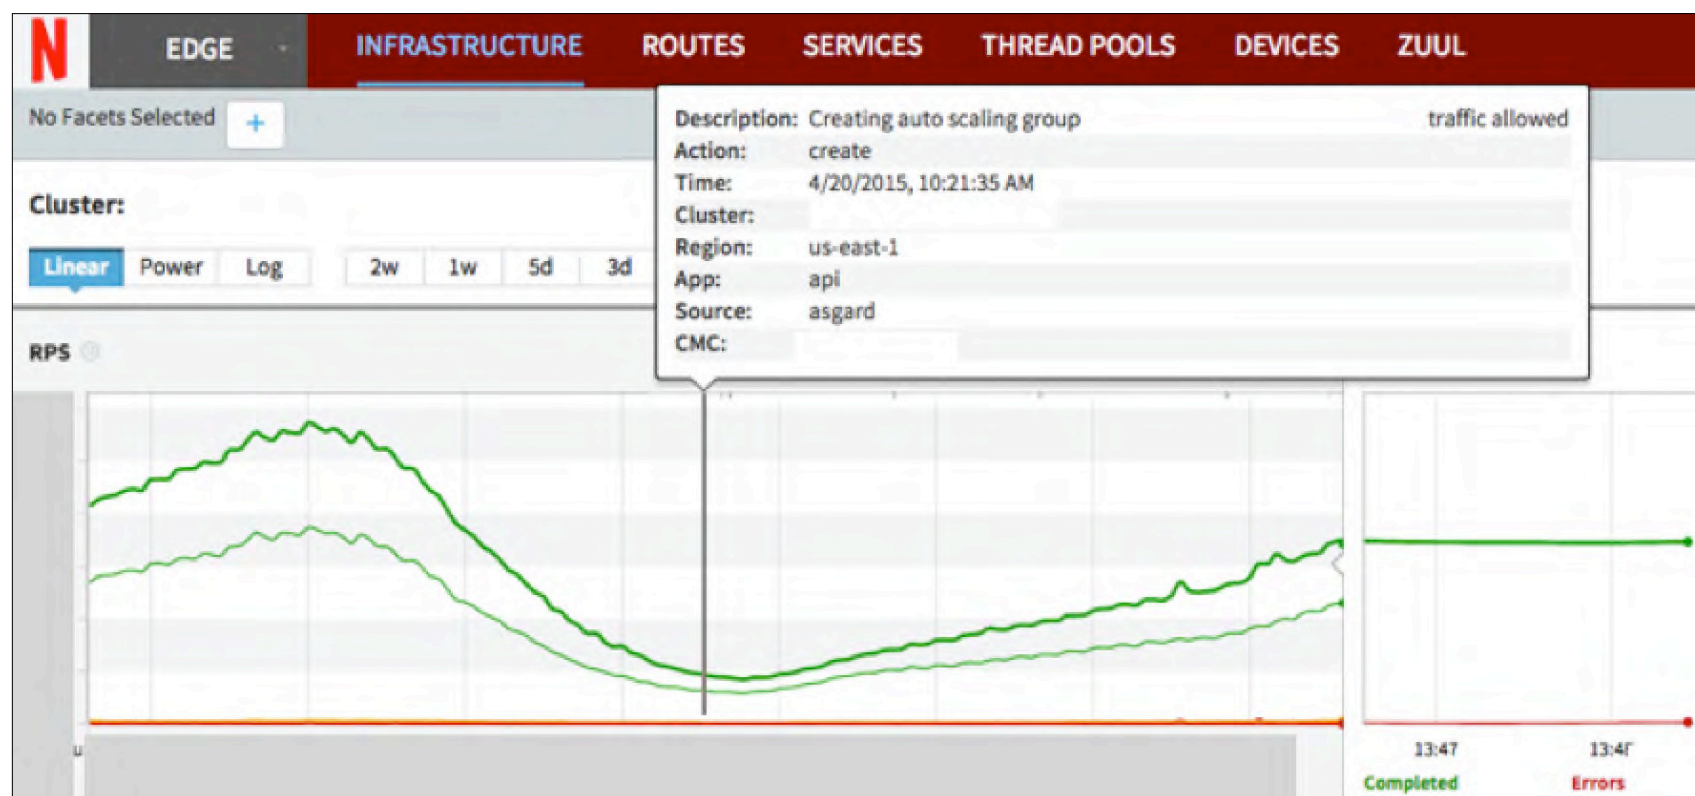
and alert configurations and decisions related to deployment schedules and rollbacks.

Powerful tools and rich visualizations equip developers with relevant information in a timely manner, providing detailed insights into code moving through the delivery pipeline as well as the health of the system in production. This section lists examples that are an integral part of the API delivery and operational process.

**Delivery insights.** The Delivery Insights tool is designed to provide the current status of all API environments, most-recent and scheduled deployments, the contents of each deployment, as well as the ability to track the flow of a given commit through the delivery pipeline (see







of operating large-scale distributed systems. Operational insight tools and dashboards need to be able to visualize massive amounts of data with relevant context. Dynamic visualizations and near-real-time data are invaluable when trying to zero in on the source of a problem.

An API operational health dashboard is also generated (see **Figure 7**). It shows a deployment event overlaid with the overall throughput trend for the API. The on-call engineers can go to this view regarding an anomaly if an alert is received, and check whether the anomaly correlated to a deployment.

This API dashboard is a sophisticated application with a Java back end and a front end written using Ember.js, D3, and RxJS. The back end connects to a stream processing system to receive data from

servers and streams it over WebSockets. These technologies enable complex data scale in real time and empower developers to detect system faults and identify opportunities for optimization.

based architecture and a heavy emphasis on data and visualization enable Netflix to move fast to scale the service, without compromising user experience. To learn more about many of the topics referenced here, visit the [Netflix Tech Blog](#).

**man** is an engineering manager at Netflix. Her  
e for software that supports Netflix' large-  
ing services running on the Amazon cloud.

## CONCLUSION

A microservices-based architecture and a heavy emphasis on automation and visualization enable Netflix to move fast while continuing to scale the service, without compromising the customer experience. To learn more about many of the tools and libraries referenced here, visit the [Netflix Tech Blog](#) or the OSS Center.

**Sangeeta Narayanan** is an engineering manager at Netflix. Her team is responsible for software that supports Netflix' large-scale, internet-facing services running on the Amazon cloud.



# How to build, use, and orchestrate Docker containers in DevOps

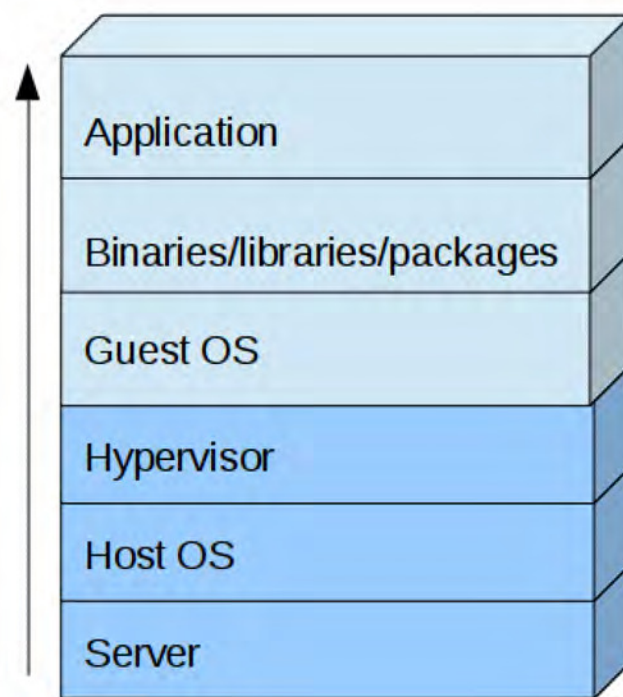
**Vagrant is a free tool** for dynamically provisioning and managing VMs.

- Traditional VMs are a great way to reduce physical hardware overhead. Nowadays, it is rarely necessary to set up physical hardware at the start of a project. It is pretty much standard to just scale out dynamically and to add further resources, in the form of more VMs. But VMs are considered to have a pretty big



footprint, particularly from a developer's perspective. Developers need to code/build/test every few minutes and won't like the repeated virtualization overhead. Let's now zoom in on what the overhead consists of and see how Docker can help.

In a VM, each virtualized application includes not only the application—which might be only a few megabytes, and the necessary binaries, libraries, and packages—but also an entire guest operating system, which might be multiple gigabytes, resulting in a comprehensive stack (see **Figure 1**).



**Figure 1.** A typical VM stack

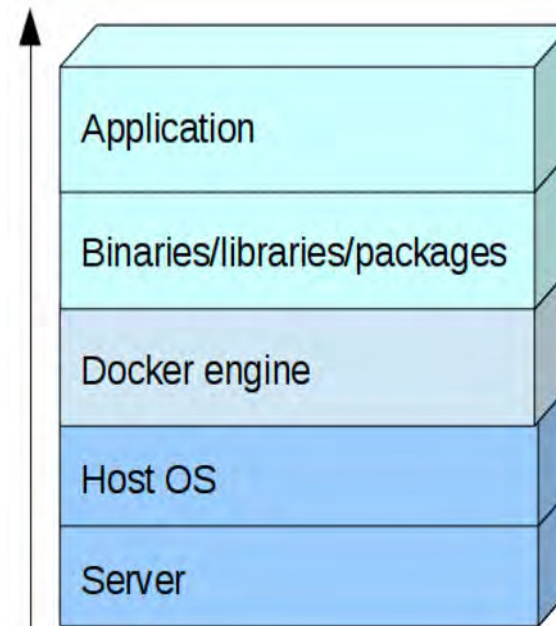
The Docker container mechanism is different. It comprises just the application and its dependencies. It runs as an isolated process in user space on the host operating system, sharing the kernel with other containers. Thus, it enjoys the resource isolation and allocation benefits of VMs but has a much smaller footprint. See **Figure 2** for an overview of what the Docker stack looks like—there is no hypervisor and there is no guest OS.

The Docker engine is based on Linux techniques such as kernel `cgroups` and namespace isolation via Linux containers (`lxc`).

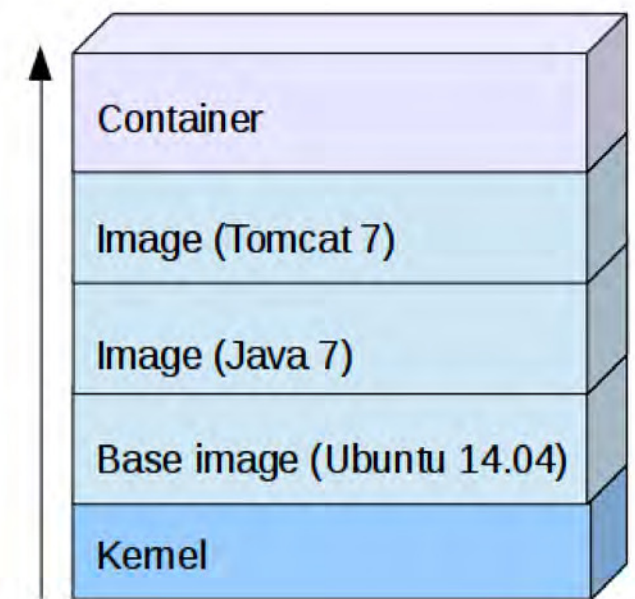
## Building a Docker Stack

Due to its lightness, Docker facilitates new behaviors for hosts, applications, and services—especially, short-lived, disposable, single services provided in a container. What was called *service-oriented architecture* (SOA) some years ago is nowadays called *microservices*.

Let's now look at a Docker use case where we slice different services, which you can also call layers or Docker containers. In practice, it is often a good choice to split containers across architectural layers or according to the frequency of change, for example, having a single Docker image for pieces that seldom change (such as middleware) and having dedicated



**Figure 2.** A typical Docker stack



**Figure 3.** The example stack

Docker images for pieces that often change (such as business applications).

My example consists of Java 7 running on a base image of Ubuntu Linux version 14.04. Because this is supposed to be a reusable image with infrastructure components, I call it the *infra* image. On top of it, I place an image running middleware—in our case, Tomcat 7 (see **Figure 3**). A business application will later be installed on this.

How is this done? As I mentioned earlier, Docker images are built from Docker definition files, named Dockerfiles, which are plain-text files located in dedicated directories. Let's have a look at the following Dockerfile:

### ■ Listing 1. (Indented lines are continuations of the previous line.)

```
FROM ubuntu:14.04
MAINTAINER Michael Huettermann
# Update Ubuntu
RUN apt-get update && apt-get -y upgrade
# Add oracle java 7 repository
RUN apt-get -y install software-properties-common
RUN add-apt-repository ppa:webupd8team/java
RUN apt-get -y update
# Accept the Oracle Java license
```





We navigate to the directory where the other Dockerfile is located. The resulting output is similar to the one above but, of course, aligned with the specific steps in that particular Dockerfile, and it closes with an individual image ID—in our case, 84b753ed5e9d.

**Docker's lightweight containers are faster** compared with classic VMs and have become popular among developers.

If I type in the `docker images` command to list the available images, I see that tomcat7 is there with a virtual size of 801.9 megabytes and the image ID just discussed.

## Running the Image

Let's now move forward from the static view (the images) to the dynamic view (the containers). The `docker run` command runs containers from images:

```
docker run -p 8080:8080 -d tomcat7
```

With this command, I created a container from our `tomcat7` container. We run Docker as a daemon (`-d`) and map ports between the host and the guest system (`-p`). I get a container identifier afterward, if everything went well. If I check afterward to see whether the newly created container runs, by using `docker ps`, I'm told it's running and that 8080 port has been mapped.

Listing all active Docker processes shows the newly created Docker container. It is surely out of scope here to go through the complete powerful Docker API, but keep in mind that you can combine a lot of shell commands in a Docker context. For example, use `docker stop $(docker ps -a -q)` to stop all containers and `docker rm $(docker ps -a -q)` to remove all containers.

The result of our Docker run is a fully functional Docker container running Java 7, Tomcat 7, and a deployed WAR file.

# Docker and Puppet

Now that we've set up the Docker container with classic Dockerfiles, you are now well prepared to learn more about how to provision Docker contain-

In our example, we now build the Tomcat 7 image, which is, as you remember, based on our first image:

```
mh@saturn:~/work/ws-git/sandbox/docker/Tomcat7$  
    docker build -t tomcat7 .
```

ers with Puppet. On the other hand, Docker itself still needs to be installed, managed, and deployed on a host. That host also needs to be managed. In turn, Docker containers might need to be orchestrated, managed, and deployed, often in conjunction with external services and tools.

The use of configuration management tools, such as Puppet, Chef [See the related article on Chef at [page 35](#). —Ed.], and Ansible, is mainstream nowadays, and they help provision target systems—including middleware, services, and business applications—in a repeatable and reliable manner.

As a result of these diverse management needs, combined with the need to manage Docker itself, I think we'll see both Docker and configuration management tools being deployed in many organizations. Indeed, there is often potential for some powerful tool chains combining containers, configuration management, continuous integration, continuous delivery, and service orchestration.

Puppet has a broad audience, because it has been on the market for years, and it enables system administrators and developers to describe the target behavior of systems in a declarative way.

## Our Example, with Puppet and Tomcat 7

Let's return to our example by setting up Tomcat 7. How does that look when using Puppet? **Figure 4** shows what the stack looks like now.

We now proceed with the Dockerfile:

```
FROM infra
MAINTAINER Michael Huettermann

RUN apt-get -y update
RUN apt-get -y install puppet librarian-puppet

RUN puppet module install puppetlabs-stdlib
RUN puppet module install puppetlabs-tomcat
RUN puppet module install puppetlabs-java

ADD site.pp /root/site.pp
RUN chmod +x /root/site.pp
ADD run.sh /root/run.sh
RUN chmod +x /root/run.sh
```

## EXPOSE 8080

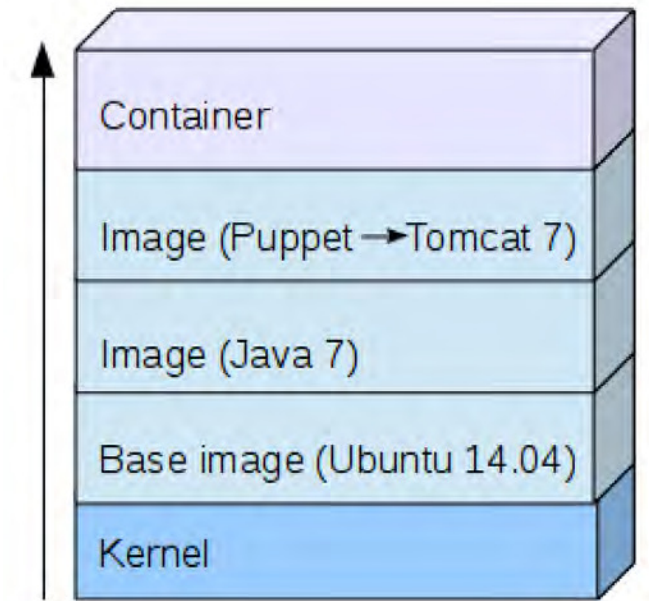
CMD ["/root/run.sh"]

Again, we derive our image from [infra](#). After some house-keeping, in particular, updating the system with new versions and installing [Librarian-puppet](#) for seamlessly installing Puppet modules from different input channels, we install some Puppet modules. Puppet, which is known for its modularity, offers many free modules to plug in for direct use. Pay special attention here to the [Puppet Tomcat module](#), which enables the installation and configuration of Tomcat. Our Dockerfile also embeds a Puppet manifest, named [site.pp](#), which is shown in **Listing 4**:

### ■ Listing 4.

```
class { 'tomcat': } ->
class { 'java': } ->
tomcat::instance { 'install':
    source_url => 'http://mirrors.ae-online.de/apache/tomcat/
        tomcat-7/v7.0.62/bin/apache-tomcat-7.0.62.tar.gz'
} ->
exec { 'run':
    command => '/opt/apache-tomcat/bin/catalina.sh start'
} ->
tomcat::war { 'all.war':
    catalina_base => '/opt/apache-tomcat',
    war_source =>
        'http://dl.bintray.com/mh/meow/all-1.0.0-GA.war',
}
```

It is out of scope here to introduce the nuts and bolts of Puppet manifests (see the Puppet documentation for that), but thanks to Puppet's declarative



### Figure 4. Our stack at present





ainers we worked on before: one that starts Tomcat and one that starts Tomcat with Puppet. To allow the host system to access both running Tomcat instances (inside the two containers, both are running on 8080), I forward the ports accordingly.

### ■ Listing 5.

```
Vagrant.configure("2") do |config|
  config.vm.define "v1" do |a|
    a.vm.provider "docker" do |d|
      d.image = "tomcat7"
      d.ports = ["8081:8080"]
      d.name = "tomcat"
    end
  end
  config.vm.define "v2" do |a|
    a.vm.provider "docker" do |d|
      d.image = "tomcat7puppet"
      d.ports = ["8082:8080"]
      d.name = "puppet"
    end
  end
end
```

That's about as straightforward an orchestration as you can find!

## On the Way to Production-Like Containers

Vagrant can serve as a Docker wrapper on systems that support Docker natively. On systems that don't, Vagrant spins up a "host VM" to run containers. You don't have to determine whether Docker is supported natively: the same configuration will work on every operating system. For that, boot2docker is used, which is a VirtualBox image of [Tiny Core Linux](#). But this tiny image seldomly reflects a production-like setup. Thus you may want to replace it with a full-fledged Linux distribution, and provision that with Docker. See **Listing 6**, where I define a Vagrant provider, which is VirtualBox with Ubuntu Trusty64 (to set up the production-like VM), and Vagrant provisioners: shell (for bootstrapping), Puppet (for setting up infrastructure), and Docker (for setting up middleware and the business application). Crisp, isn't it?

■ **Listing 6.**

```
Vagrant.configure(2) do |config|
  config.vm.define "with-docker"
  config.vm.provider "virtualbox" do |v|
    v.name = "_withDocker"
  end
  config.vm.box = "ubuntu/trusty64"
  config.vm.box_check_update = false
  config.vm.provision :shell, path: "bootstrap.sh"
  config.vm.provision "puppet" do |puppet|
    puppet.manifests_path = "manifests"
    puppet.manifest_file = "default.pp"
  end
  config.vm.provision "docker" do |d|
    d.build_image "/vagrant/docker/Tomcat7PuppetFull",
      args: "-t tomcat7dockerfromvagrant"
    d.run "tomcat7dockerfromvagrant",
      args: "-p 8080:8080"
  end
end
```

## Conclusion

Docker's lightweight containers are faster compared with classic VMs and have become popular among developers and as part of CD and DevOps initiatives. If your purpose is isolation, Docker is an excellent choice.

Vagrant is a VM manager that enables you to script configurations of individual VMs as well as do the provisioning. However, it is still a VM dependent on VirtualBox (or another VM manager) with relatively large overhead. It requires you to have a hard drive file that can be huge, it takes a lot of RAM, and performance can be suboptimal. Docker uses kernel cgroups and namespace isolation via lxc. This means that you are using the same kernel as the host and the same file system. Vagrant is a level above Docker in terms of abstraction, so they are not really comparable.

Configuration management tools such as Puppet are widely used for provisioning target environments. Reusing existing Puppet-based solutions is easy with Docker. You can also slice your solution, so the infrastructure is provisioned with Puppet; the middleware, the business application itself, or both are provisioned with Docker; and Docker is wrapped by Vagrant. With this range of tools, you can do what's best for your scenario. </article>



**BY ALAN FRYER**

This article demonstrates how to implement CD using open source tools such as Git, Maven, Nexus, Jenkins, and

The first step is modeling infrastructure with code, which I will do using the open source tool, [Chef](#). This step will enable much of the automation of the other tools. In this article, I'll be modeling the configuration illustrated in **Figure 1**. As you can see, the elements I just described are present, along with a Chef server, which I'll describe in a moment. The additional machines, on which a Chef client is installed, are referred to as *nodes*.









We have now created a mechanism that automates the deployment of applications to multiple nodes from a single command-line call.

When the file is saved and the editor is closed, the data bag is created on the server in an encrypted format. An encrypted data bag must be created to contain credentials for each environment. The password for the management user `mgm_pwd` is read from the data bag using the recipe in **Listing 4A**.

**Install Postgres module.** The `installpostgresmodule` recipe in **Listing 5** installs the Postgres JDBC driver in the JBoss `modules` folder. The folder `$JBOSS_HOME/modules/postgres/main` is created using the `directory` resource and copies the cookbook file `module.xml` (**Listing 6**) to this folder using the `cookbook file` resource.

**Start JBoss Standalone instance.** The `startstandalone` recipe in **Listing 7** starts the JBoss Standalone instance by using the `execute` resource to run the shell script `$JBOSS_HOME/bin/standalone.sh` (**Listing 7A**) in `nohup` mode.

The management address and JBoss server is bound to the DNS name of the server by setting the options `-b` and `-Djboss.bind.address.management` with the value obtained from the node attribute `node['fqdn']`.

**Configure JBoss Standalone instance.** The `configurestandalone` recipe in **Listing 8** configures the Postgres datasource, downloads the application WAR file from Nexus, and deploys it to the JBoss server. The configuration and deployment are performed by executing a JBoss CLI script against the running server. The CLI script (**Listing 9**) is defined as a cookbook template, which is an embedded Ruby template that is used to generate the CLI script based on the variables and logic contained within the template. The template is generated using the following resource definition, which passes in the variables read from the node:

```
template cli_script do
  source "configure-standalone.cli.erb"
  variables(
    ds_jndi_name: "postgresDS",
    ds_pool_name: "postgresDS",
    db_host: node['jboss']['database_hostname'],
    db_port: node['jboss']['database port'],
```

```

    db_name: node['jboss']['database_name'],
    db_user: node['jboss']['database_user'],
    db_password: creds["postgres"],
    jboss_install: jboss_install,
    artifact: "#{artifact_id}.war",
    version: version
  )
  owner process_user
  group process_group
end

```

The `execute` resource then runs the shell script `$JBOSS_HOME/bin/jboss-cli.sh` with the option `--file=configure-standalone.cli` to run the CLI script generated from the template.

## INTEGRATING CHEF AND JENKINS

The Chef client and `knife` tool need to be configured on the server running Jenkins. This step requires creating the user `jenkins` on the Chef server using the management console, and linking it to an organization—in this case, C2B2. A private key for the user is generated. It is saved to the `knife` tool's configuration folder `/home/jenkins/.chef`. The file `knife.rb` configures the connection to the Chef server, together with the location of the private key used to authenticate the connection and location of the cookbook's folder in the local Chef repository (**Listing 10**).

To execute a Chef run, a Jenkins Freestyle Project is used. The build checks out the Chef code from a specific branch in a Git repository to the build's workspace, relying on `CHEF_ENVIRONMENT` and `CHEF_ROLE`. The checked-out code includes the shell script in **Listing 11**, which is executed as part of the build.

The shell script uploads the checked-out version of the cookbook from the Git repository branch configured in the build, which executes the `knife ssh` command, running the Chef client on the nodes matching the environment and the role passed in to the script.







# Functional Programming in Java: Using Collections

**I**n this two-part article, I demonstrate how to use Java 8 lambda expressions to take advantage of the functional style of programming to create more-expressive and concise code with less mutability and fewer errors.

Collections of numbers, strings, and objects are used so commonly in Java that removing even a small amount of ceremony from coding them can reduce code clutter greatly. In this article, I explore the use of lambda expressions to manipulate collections. I use them to iterate collections, extract elements from them, and easily concatenate their elements.

## Iterating Through a List

Iterating through a list is a basic operation on a collection, but over the years it's gone through a few significant changes. We'll begin with the old and evolve an example—enumerating a list of names—to the elegant style.

We can easily create an immutable collection of a list of names with the following code:

Here's the habitual, but not so desirable, way to iterate and print each of the elements.

I call this style the *self-inflicted wound pattern*—it’s verbose and error prone. We have to stop and wonder, “Is it `i <` or `i <=`?” This is useful only if we need to manipulate elements at a particular index in the collection, but even then, we can opt to use a functional style that favors immutability, as we’ll discuss soon.

Java also offers a construct that is a bit more civilized than the good old `for` loop.

Under the hood this form of iteration uses the `Iterator` interface and calls into its `hasNext()` and `next()` methods.

Both these versions are *external iterators*, which mix *how* we do it with *what* we'd like to achieve. We explicitly control the iteration with them, indicating where to start and where to end; the second version does that under the hood using the `Iterator` methods. With explicit control, the `break` and



`continue` statements can also help manage the iteration's flow of control.

The second construct has less ceremony than the first. Its style is better than the first if we don't intend to modify the collection at a particular index. Both of these styles, however, are imperative and we can dispense with them in modern Java by using a functional approach.

There are quite a few reasons to favor the change to the functional style:

- The **for** loops are inherently sequential and are quite difficult to parallelize.
- Such loops are nonpolymorphic; we get exactly what we ask for. We passed the collection to **for** instead of invoking a method (a polymorphic operation) on the collection to perform the task.
- At the design level, the code fails the “Tell, don’t ask” principle. We ask for a specific iteration to be performed instead of leaving the details of the iteration to underlying libraries.

It's time to trade in the old imperative style for the more elegant functional-style version of *internal iteration*. With an internal iteration we will fully turn over most of the *hows* to the underlying library so we can focus on the essential *whats*. The underlying function will take care of managing the iteration. Let's use an internal iterator to enumerate the names.

The `Iterable` interface has been enhanced in JDK 8 with a special method named `forEach()`, which accepts a parameter of type `Consumer`. As the name indicates, an instance of `Consumer` will consume, through its `accept()` method, what's given to it. Let's use the `forEach()` method with the all-too-familiar anonymous inner class syntax.

```
friends.forEach(new Consumer<String>() {
    public void accept(final String name) {
        System.out.println(name);
    }
});
```

Here, we have invoked `forEach()` on the `friends` collection and passed an anonymous instance of `Consumer` to it. The `forEach()` method will invoke the `accept()` method of the given `Consumer` for each element in the collection and perform a specified action. In this example, it merely prints the given value, which is a name.

Let's look at the output from this version, which is the same as the output from the two previous versions:

Brian  
Nate  
Neal  
Raju  
Sara  
Scott

We changed just one thing: we traded in the old `for` loop for the new internal iterator `forEach()`. As for the benefit, we went from specifying how to iterate to focusing on what we want to do for each element. The bad news is the code looks a lot more verbose—so much that it can drain away any excitement about the new style of programming. Thankfully, we can fix that quickly; this is where lambda expressions and the new compiler magic come in. Let's make one change again, replacing the anonymous inner class with a lambda expression.

```
friends.forEach((final String name) ->
    System.out.println(name));
```

That's a lot better. We look at less code, but watch closely to see what's in there. The `forEach()` is a higher-order function that accepts a lambda expression or block of code to execute in the context of each element in the list. The variable `name` is bound to each element of the collection during the call. The underlying library takes control of how any lambda expressions are evaluated. It can decide to perform them lazily, in any order, and exploit parallelism as it sees fit.

This version produces the same output as the previous versions. The internal-iterator version is more concise than the other ones. In addition, when we use it we're able to focus our attention on what we want to achieve for each element rather than how to sequence through the iteration—it's declarative.

This version has a limitation, however. Once the `forEach` method starts, unlike in the other two versions, we can't break out of the iteration. (There are facilities to handle this limitation.) As a consequence, this style is useful

**The Java compiler treats single-parameter lambda expressions as special.** We can leave off the parentheses around the parameter if the parameter's type can be inferred.









From the output, it's clear that the method picked up the proper number of elements from the input collection.

Found 2 names

The `filter()` method returns an iterator just like the `map()` method does, but the similarity ends there. Whereas the `map()` method returns a collection of the same size as the input collection, the `filter()` method might not. It might yield a result collection with a number of elements ranging from zero to the maximum number of elements in the input collection. However, unlike `map()`, the elements in the result collection that `filter()` returned are a subset of the elements in the input collection.

## Conclusion

The conciseness we've achieved by using lambda expressions so far is nice, but code duplication might sneak in quickly if we're not careful. I'll address this concern in the second half of this article in the upcoming issue. **</article>**

*This article was adapted from [Functional Programming in Java: Harnessing the Power of Java 8 Lambda Expressions](#) with kind permission from the publisher, The Pragmatic Bookshelf.*

**LEARN MORE**

- [Video of the author lecturing on this topic](#)
- [What is higher-order programming?](#)
- [Why functional programming is on the rise](#)





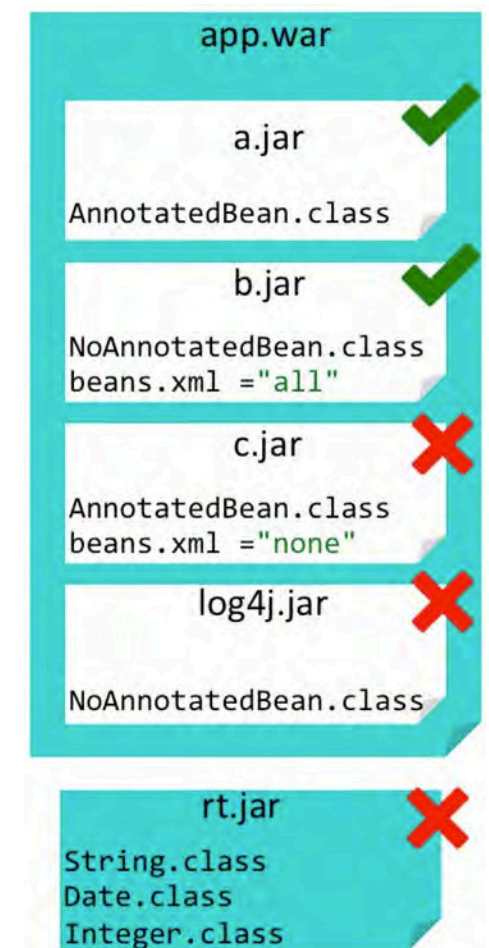
# Contexts and Dependency Injection: the New Java EE Toolbox

This series of four articles attempts to demystify Contexts and Dependency Injection (CDI). In the [previous article](#), I discussed what strong typing means in dependency injection. Here, I explain how to integrate third-party frameworks. The next article will focus on how to achieve loose coupling with interceptors, decorators, and events. The final article will cover the integration of CDI within Java EE.

In the previous article, I examined how CDI can inject beans into other beans. The problem we now face is that when we want to integrate third-party frameworks, their classes are not discovered by CDI and so they can't be injected. Thanks to producers, however, CDI enables us to turn any Java class into a CDI bean so it can be managed by the CDI container. Once a class is managed, we can use qualifiers and alternatives as we saw in the previous article. Thanks to disposers, a produced bean can be cleaned up easily.

When an application starts, the CDI container performs *bean discovery*. This is the process in which the container searches for CDI beans in all bean archives that are declared in the application classpath. During bean discovery, the CDI container detects

When we talk about a *CDI bean*, we simply mean a Java class that follows certain patterns and is managed by a CDI container. A CDI bean must be a concrete class with a default constructor; it can have



### Figure 1. What constitutes a bean archive



optional qualifiers, scopes, an expression language name, or a set of interceptor bindings. In fact, with very few exceptions, potentially every Java class that has a default constructor can be a CDI bean, if packaged in a bean archive.

## Bean Archives

CDI doesn't define any particular deployment archive. We can package CDI beans in JARs or WARs. However, CDI discovers only beans that are packaged following certain rules—those in a *bean archive*. A bean archive is any archive that packages Java classes annotated with CDI annotations. For example, if a class contains a CDI annotation (qualifier, interceptor binding, and so on), that class is said to be an **AnnotatedBean**; if it appears in **a.jar**, then **a.jar** is said to be a bean archive.

The other possibility is to use a `beans.xml` deployment descriptor in a JAR with a bean discovery mode set to `all`. In **Figure 1**, `b.jar` packages no annotated beans, but because it contains a `beans.xml` file, it is a bean archive.

Any other packaging is said to be a plain archive. For example, if the bean discovery mode in `beans.xml` is set to `none` (meaning that we tell CDI not to discover any beans in this archive), then `c.jar` is not treated as a bean archive. This means that even if `c.jar` packages CDI beans, they will not be discovered and not managed by CDI. The other common case is when the archive does not contain any annotated beans and has no `beans.xml` descriptor. In **Figure 1**, `log4j.jar` and `rt.jar` do not contain any of those, so they are not bean archives.

When we package a web application in a WAR file, it is common to bundle third-party frameworks that are not bean archives. For example, the `app.war` web application needs Log4j, a third-party logging framework, and of course, Java APIs such as `String`, `Date`, and `Integer`, which are packaged in the `rt.jar` file. This `rt.jar` file contains the Java runtime environment (JRE) classes and doesn't have a `beans.xml` deployment descriptor. So, these classes cannot be discovered or managed by CDI, and therefore they cannot be injected. To be injected, they first need to be *produced*.

## Producers

Basically, a *producer* provides CDI capabilities to a plain old Java object (POJO) or to any datatype. Why would we want to do that? The answer is that there are plenty of cases where we need additional control. What if we

need to decide at runtime which implementation of a datatype or logger to instantiate and inject?

That's where producers are beneficial: they enable third-party frameworks to be used with CDI by exposing their objects as CDI beans. To do this, declare a field or a method to be produced by annotating it with `@Produces`.

**Producer fields.** Let's look at an example that illustrates producers. In **Listing 1**, the `BookService` class has a method to create a book object given a title. This method also generates an ISBN number by concatenating a prefix, a random number, and a postfix. These attributes are part of the service and, as you can see, are not initialized.

### ■ Listing 1.

```
public class BookService {

    @Inject
    private String prefix;
    @Inject
    private int random;
    @Inject
    private long postfix;

    public Book createBook(String title) {
        String number = prefix + "-" +
            random + "-" + postfix;
        return new Book(title, number);
    }
}
```

The goal is to ask CDI to inject the values of these attributes. As mentioned earlier, **String**, **int**, and primitive datatypes are part of the JRE, and therefore cannot be injected. The only way for this to work is to produce them.

To do this, we just take any class in our application, such as `NumberProducer` (**Listing 2**), but it could have been called anything else. We declare some attributes, initialize them, and ask CDI to produce them with the `@Produces` annotation. This means that all the produced datatypes can now be injected with `@Inject` anywhere in our application. Notice that CDI does not use the name of the class and the name of

the attributes to do the binding. Here, the prefix attribute is called **pre**, but it could have been called anything else. CDI doesn't use the name but the type. CDI knows that one attribute is of type **String**, another **long**, and the other **int**. That's why CDI is said to be strongly typed.

### ■ Listing 2.

```
public class NumberProducer {

    @Produces
    private String pre = "7";

    @Produces
    private int random =
        Math.abs(new Random().nextInt());

    @Produces
    private long post =
        Math.abs(new Random().nextLong());

}
```

Going back to our `BookService` class (**Listing 1**), the important information here for CDI is the datatype. Being strongly typed, CDI will inject the right value to the right datatype. This eliminates lookup using string-based names or XML for wiring so that the compiler will detect any errors. But as you might have guessed, datatypes are not enough. What if we need to inject another attribute of type `long` with a completely different value—for example, milliseconds? That’s when we use qualifiers on producers.

**Injecting qualified produced datatypes.** If we need different values to be injected for the same datatype, we use qualifiers. In **Listing 3**, we change the `BookService` to add a qualifier on `postfix` and ask CDI to inject a 13-digit postfix. With a different qualifier, `@CurrentTime`, we tell CDI to inject the current time in milliseconds.

### ■ Listing 3.

```
public class BookService {  
  
    @Inject  
    private String prefix;
```

```
@Inject
private int random;

@Inject
@ThirteenDigits
private long postfix;

@Inject
@CurrentTime
private long millis;

public Book createBook(String title) {
    String number =
        prefix + "-" + random + "-" +
        postfix + "-" + millis;
    return new Book(title, number);
}
```

For this to work, we need to go back to the producer class and add the right qualifiers on the producers (see **Listing 4**). We can read this code as follows: produce a `String` called `@Default`, produce an `int` called `@Default`, produce a `long` called `@ThirteenDigits`, and produce another `long` called `@CurrentTime`. Being produced, they are all managed, and therefore injectable.

■ **Listing 4.**

```
public class NumberProducer {

    @Produces
    private String pre = "7";

    @Produces
    private int random =
        Math.abs(new Random().nextInt());

    @Produces
    @ThirteenDigits
```





```
        return file;
    }
}
```

The way to use this produced file is easy. Take the `FileService` bean defined in **Listing 7**. It has a single method that writes some UTF-8 text to a given file. Instead of going through the long process of creating the file, it just injects it. A producer has already created the injected file, and it just needs to be injected so it can be used. Our system could then produce several files, just by using qualifiers. For example, we could create temporary files by qualifying the injection point with `@Temp`, or we could inject files used to store large binary objects with `@Lob`. As you can see, in this example, producer methods are very valuable. But producers are themselves CDI beans in their own right. So producers can use injection and other CDI features.

■ **Listing 7.**

```
public class FileService {

    @Inject
    private Path file;

    public void write(String text) throws Exception {
        Files.write(file, text.getBytes("utf-8"));
    }

}
```

**Producers are CDI beans.** Let us take the `FileProducer` bean and slightly change the file creation algorithm (**Listing 8**). This time, the directory to be created is injected. Being a managed bean, `FileProducer` can inject values like any other bean. In this example, it injects the root directory using the `@Root` qualifier (which has been produced somewhere else). The filename also changes by adding the number of milliseconds to the name. This number of milliseconds is injected and represents the current time. The file is created, returned, and produced, and it can now be injected.

■ **Listing 8.**

```
public class FileProducer {
```

```
@Inject
@Root
private Path directory;

@Inject
@CurrentTime
private long millis;

@Produces
public Path produceFile() throws IOException {
    Set<PosixFilePermission> perms =
        PosixFilePermissions.fromString("rwxr-x---");
    FileAttribute<Set<PosixFilePermission>> attr =
        PosixFilePermissions.asFileAttribute(perms);
    if (Files.notExists(directory))
        Files.createDirectory(directory, attr);
    Path file =
        directory.resolve("myfile-" + millis + ".txt");
    if (Files.notExists(file))
        Files.createFile(file, attr);
    logger.info("File {} created", file);
    return file;
}
```

## InjectionPoint API

Until now, the produced attributes and returned values that we have seen did not need any information about where they were injected. But there are certain cases where produced objects need to know something about the injection point into which they are injected. This can be a way of configuring or changing the producers' behavior depending on the injection point. CDI has an **InjectionPoint** API that provides access to metadata about an injection point. Thus we can create a producer method that has an **InjectionPoint** as a parameter.

Let's take for example the creation of a logger with the external Log4j library. Here, the `BookService` class creates a logger and sets its name to `BookService.class`.

```
public class BookService {  
  
    Logger log = getLogger(BookService.class);  
  
    // ...  
}
```

If instead we wanted to inject the logger, we would need to produce it. For that, we take a separate **LoggingProducer** class, add a producer method that returns the logger, and inject it. It's as simple as that. But what if we need to reuse this logger in other classes? All the loggers will be named **BookService**, and that's not what we want. If **BookService** needs a logger called **BookService**, that means that **FileService** needs its own logger, as well as **ItemService**. The only parameter that changes here is the name of the logger.

```
public class FileService {  
  
    Logger log = getLogger(FileService.class);  
  
    // ...  
}
```

How would we produce a logger that needs to know the class name of the injection point? The answer is by using the `InjectionPoint` API. This API provides access to metadata about an injection point, and in our case, it is used to configure the logger. In **Listing 9**, the `LoggingProducer` class produces a parameterized logger, thanks to `injectionPoint.getMember().getDeclaringClass().getName()`. This will get the injection point class name (`BookService`, `FileService`, or `ItemService`) and produce a logger per class. The `InjectionPoint` API has several methods to return the bean type of the injection point, its qualifiers, or the object itself.

### ■ Listing 9.

```
public class LoggingProducer {  
    @Produces
```

```
public Logger produceLogger(
    InjectionPoint injectionPoint) {
    return LogManager
        .getLogger(injectionPoint.getMember()
            .getDeclaringClass().getName());
}
```

Using this produced parameterized logger is straightforward: we just inject it and use it as usual. The logger's category class name will then be set automatically, thanks to the **InjectionPoint** API. We can now do the same thing for **FileService** and **ItemService**. The code is strictly the same, but each bean will be injected with its own parameterized logger. Writing that simple parameterized producer has saved us time. We don't need to retrieve the class name to set it on a logger in every class where we want to use it. For example,

```
public class BookService {

    @Inject
    private Logger logger;

    // ...

}

public class FileService {

    @Inject
    private Logger logger;

    // ...

}
```

## Disposers

In the previous examples, I used producers to create datatypes or POJOs so they could be managed by CDI. Producers act as factories: they create manageable objects. We have created objects and didn't destroy or close them, because we didn't need to. But some producer methods can return objects that require explicit destruction, such as a JDBC connection, a Java

Message Service session, or an entity manager. For destruction, CDI uses *disposers*. A disposer method enables the application to perform the customized cleanup of an object returned by a producer method. Let's look at an example using JDBC.

**Producing a JDBC connection.** In Listing 10, the `JDBCConnectionProducer` class has a `createConnection` method that creates a JDBC connection. This method takes a Derby database JDBC driver, creates a connection with a specific URL, and returns an opened JDBC connection. This method is annotated with `@Produces`, meaning that the connection can now be injected.

■ **Listing 10.**

```
public class JDBCConnectionProducer {

    @Inject
    private Logger logger;

    @Produces
    private Connection createConnection()
        throws Exception {
        Class.forName(
            "org.apache.derby.jdbc.EmbeddedDriver")
            .newInstance();
        Connection conn =
            DriverManager.getConnection(
                "jdbc:derby:memory:myDB;create=true",
                "APP", "APP");
        logger.info("Connection created");
        return conn;
    }
}
```

Now let's take a `JDBCPingService` (see **Listing 11**) whose job is just to ping a database to see if it's up and running. The `ping` method executes a SQL statement on a JDBC connection—that is, the produced connection that we can now inject. This code is fine, but we need to close the JDBC connection. One solution is to do it in the `JDBCPingService` by calling the `close` method. But this doesn't look right: if one producer has created this

connection, it should be its job to close it. That's the role of disposers.

■ **Listing 11.**

```
public class JDBCPingService {

    @Inject
    private Logger logger;

    @Inject
    private Connection conn;

    public void ping() throws SQLException {
        conn.createStatement().executeQuery(
            "SELECT 1 FROM SYSIBM.SYSDUMMY1");
        logger.info("Ping....");
        conn.close();
    }
}
```

**Disposing a JDBC connection.** So we go back to our `JDBCConnectionProducer` class and add a `closeConnection` method (see **Listing 12**). The role of this method is to terminate the JDBC connection. Notice that both methods here are `private` and can't be invoked by other classes; only CDI can manage them.

To be able to automatically invoke the `closeConnection` method, it has to be annotated with `@Disposes`. Destruction is performed by a matching disposer method. Each disposer method must have exactly one disposed parameter of the same type as the corresponding producer method returned type—here, `Connection`. The disposer method, `closeConnection`, is called automatically when the client context ends (`@ApplicationScoped`, `@SessionScoped`, `@RequestScoped`, and so on).

■ **Listing 12.**

```
public class JDBCConnectionProducer {

    @Inject
    private Logger logger;
```



```
@Produces
private Connection createConnection()
    throws Exception {
    Class.forName(
        "org.apache.derby.jdbc.EmbeddedDriver")
        .newInstance();
    Connection conn =
        DriverManager.getConnection(
            "jdbc:derby:memory:myDB");
    logger.info("Connection created");
    return conn;
}

private void
    closeConnection(@Disposes Connection conn)
        throws SQLException {
    conn.close();
    logger.info("Connection closed");
}
}
```

The way to use a produced and disposed connection is as expected. The `JDBCPingService` in **Listing 11** changes just by getting rid of the `conn.close()` statement. It injects the produced JDBC connection with `@Inject`, and the `ping` method uses it to ping the Derby user database. The `JDBCPingService` class doesn't need to deal with all the technical plumbing of creating and closing the JDBC connection or exception handling. Producers and disposers are a tidy way of creating and closing resources.

**Qualifying producer and disposer.** *Qualifiers* are a type-safe way to distinguish between several beans and injection points, and can also be used on producers. For example, let's say I have two databases: a user database and a purchase order database.

To produce and dispose connections for these two databases, we just qualify the producer and disposer (**Listing 13**) with the `@UserDatabase` or `@PurchaseOrderDatabase` qualifier.

■ **Listing 13.** (Full listing available in downloads.)

```
public class JDBCConnectionProducer {

    @Inject
    private Logger logger;

    @Produces
    @UserDatabase
    private Connection
        createUserConnection() throws Exception {
        Class.forName(
    . . .
    @Produces
    @PurchaseOrderDatabase
    private Connection createPOConnection()
        throws Exception {
        Class.forName(
    . . .
}
```

For a given producer, there can be only one disposer method that matches the bean type and the qualifier. Also, notice that a disposer method must be declared within the same class as the producer. In **Listing 14** ([download](#) only), I have enriched the `JDBCPingService` class so it can ping both user and purchase order databases.

## Alternatives and Producers

The benefit of *alternatives* is to change the behavior of an application at deployment time, just by notation in the `beans.xml` file. Alternatives can be applied to beans, as well as producers, just by using the `@Alternative` annotation.

Let's examine an example. In **Listing 15**, the **PurchaseOrderService** injects a value-added tax (VAT) rate and a discount rate. Both attributes of type **Float** are used to create a purchase order. Given a subtotal, the **PurchaseOrderService** sets both rates and calculates the total of the purchase order. As you might have guessed by now, both the VAT rate and the discount rate are produced in a separate class. Thanks to the **@Vat** and **@Discount** qualifiers, both fields of the same datatype, **Float**, can be produced without any ambiguity.

■ **Listing 15.**

```
public class PurchaseOrderService {

    @Inject
    @Vat
    private Float vatRate;

    @Inject
    @Discount
    private Float discountRate;

    public PurchaseOrder
        createPurchaseOrder(Float subTotal) {
        PurchaseOrder order =
            new PurchaseOrder(subTotal);
        order.setVatRate(vatRate);
        order.setDiscountRate(discountRate);
        order.setTotal(subTotal + (
            subTotal * vatRate) - (
            subTotal * discountRate));
        return order;
    }
}
```

But now let's say I deploy the application in a different environment or country, and I need to change the values of both rates. The goal is to produce another VAT rate and make sure it is annotated with `@Alternative`. We do the same for the discount rate. We end up with a single class, `NumberProducer` (see **Listing 16**), producing a VAT rate (5.5 percent) and an alternative VAT rate (7.8 percent) within the same bean.

■ **Listing 16.**

```
public class NumberProducer {  
  
    @Produces  
    @Vat  
    private Float vatRate = 0.055f;  
}
```

```
@Produces
@Vat
@Alternative
private Float altVatRate = 0.078f;

@Produces
@Discount
private Float discountRate = 0.0225f;

@Produces
@Discount
@Alternative
private Float discountRateAlt = 0.125f;
}
```

To enable these produced `@Alternative` values, we just need to add the `NumberProducer` to the `beans.xml` deployment descriptor (**Listing 17**). The `PurchaseOrderService` in **Listing 15** doesn't change. Depending on whether the alternative is activated or not, the VAT rate of 5.5 percent will be injected, or the alternative one of 7.8 percent.


■ **Listing 17.**

```
<beans xmlns="http://xmlns.jcp.org/xml/ns/javaee"
    . . .
    version="1.1" bean-discovery-mode="all">
  <alternatives>
    <class>org.foo.bar.NumberProducer</class>
  </alternatives>
</beans>
```

## Conclusion


CDI makes it very easy to convert third-party classes into injectable objects using the `beans.xml` file and a variety of annotations. The annotations are remarkably simple, and they support a wide range of programming needs. As I've shown here, CDI creates one of the easiest ways to build apps using loosely coupled third-party components. **</article>**

Written by leading Java experts, Oracle Press books offer the most definitive, complete, and up-to-date coverage of Java available.




**Java EE 7: The Big Picture**  
*Danny Coward*

Master the code, applications, and frameworks that power Java Platform, Enterprise Edition 7.



**OCA/OCP Java SE 7  
Programmer I & II Study Guide**  
*Kathy Sierra, Bert Bates*

This self-study tool offers full coverage of OCA/OCP Exams 1Z0-803 and 1Z0-804. Electronic content includes 500+ practice exam questions.



**Mastering JavaFX 8 Controls**  
*Hendrik Ebbers*


Create custom Java FX 8 controls and deliver state-of-the-art applications with visually stunning UIs.



**Java EE Applications  
on Oracle Java Cloud**  
*Harshad Oak*

Build highly available, scalable, secure, distributed applications on Oracle Java Cloud.





MICHAEL HEINRICHS



## Part 2

# Inside the CPU: the Unexpected Effects of Instruction Execution

## How false sharing and branch misprediction can have unwanted effects on your code's performance

The first article in this series about the effects of modern chip design on programming focused on the memory system, specifically the cache hierarchy. In this article, Part 2, I extend our investigation to the issues of multithreaded access. In Part 3, I will look at the internal workings of the CPU itself.

## False Sharing

In our first experiment, we will see a surprising effect that can occur when memory is accessed concurrently from different threads. The code in **Listing 1** contains an array of 17 integers and four methods that modify different elements of the array. The experiment consists of two test runs in which we run two of the methods concurrently on different threads. In the first test run, we will execute the methods `modifyFarA()` and `modifyFarB()`, which modify two array elements that are 16 elements apart. In the second test run, we will execute the methods `modifyNearA()` and `modifyNearB()`, which modify two adjacent array elements. The only difference is how far apart the modified array elements are. How does that affect performance?

### ■ Listing 1.

```
public final int[] array = new int[17];
```

```
@Benchmark
@Group("near")
public void modifyNearA() {
```

```

        array[0]++;
    }

    @Benchmark
    @Group("near")
    public void modifyNearB() {
        array[1]++;
    }

    @Benchmark
    @Group("far")
    public void modifyFarA() {
        array[0]++;
    }

    @Benchmark
    @Group("far")
    public void modifyFarB() {
        array[16]++;
    }
}

```

On my laptop, a method call in the first test run takes about 3.6 ns, while in the second test run it takes 4.5 ns. In other words, if the array elements are farther apart, modifications are 25 percent faster. How can that be?

There are actually two puzzles to solve here. Why does it matter how far apart the elements are? And why do these

methods interfere with each other at all, even though they access different variables?

Watchful readers will notice the similarity to the initial example in the first article of this series. If elements of an integer array are 16 elements or more apart, they will be located in different cache lines. If they are closer to each other, chances are great that they will end up in the same cache line. Is the observed variation in behavior related to cache lines?

Indeed, it is. The memory system has no understanding of Java. It does not know about Java arrays and Java variables. Its smallest unit is a cache line. If two threads modify the same cache line, they interfere with each other and it does not matter that we modified two different variables in the source code.

Why do two threads that modify the same cache line slow each other down? We have to fill out our model of the cache hierarchy to understand this effect. When two threads run in parallel long enough, they end up on different CPU cores. CPU cores do not share an L1 cache—each core has its own. (Note: How the L2 and L3 caches are shared between cores depends on the architecture. For example, a typical setup is that L2 cache is shared between adjacent cores while L3 cache is shared between all cores.)

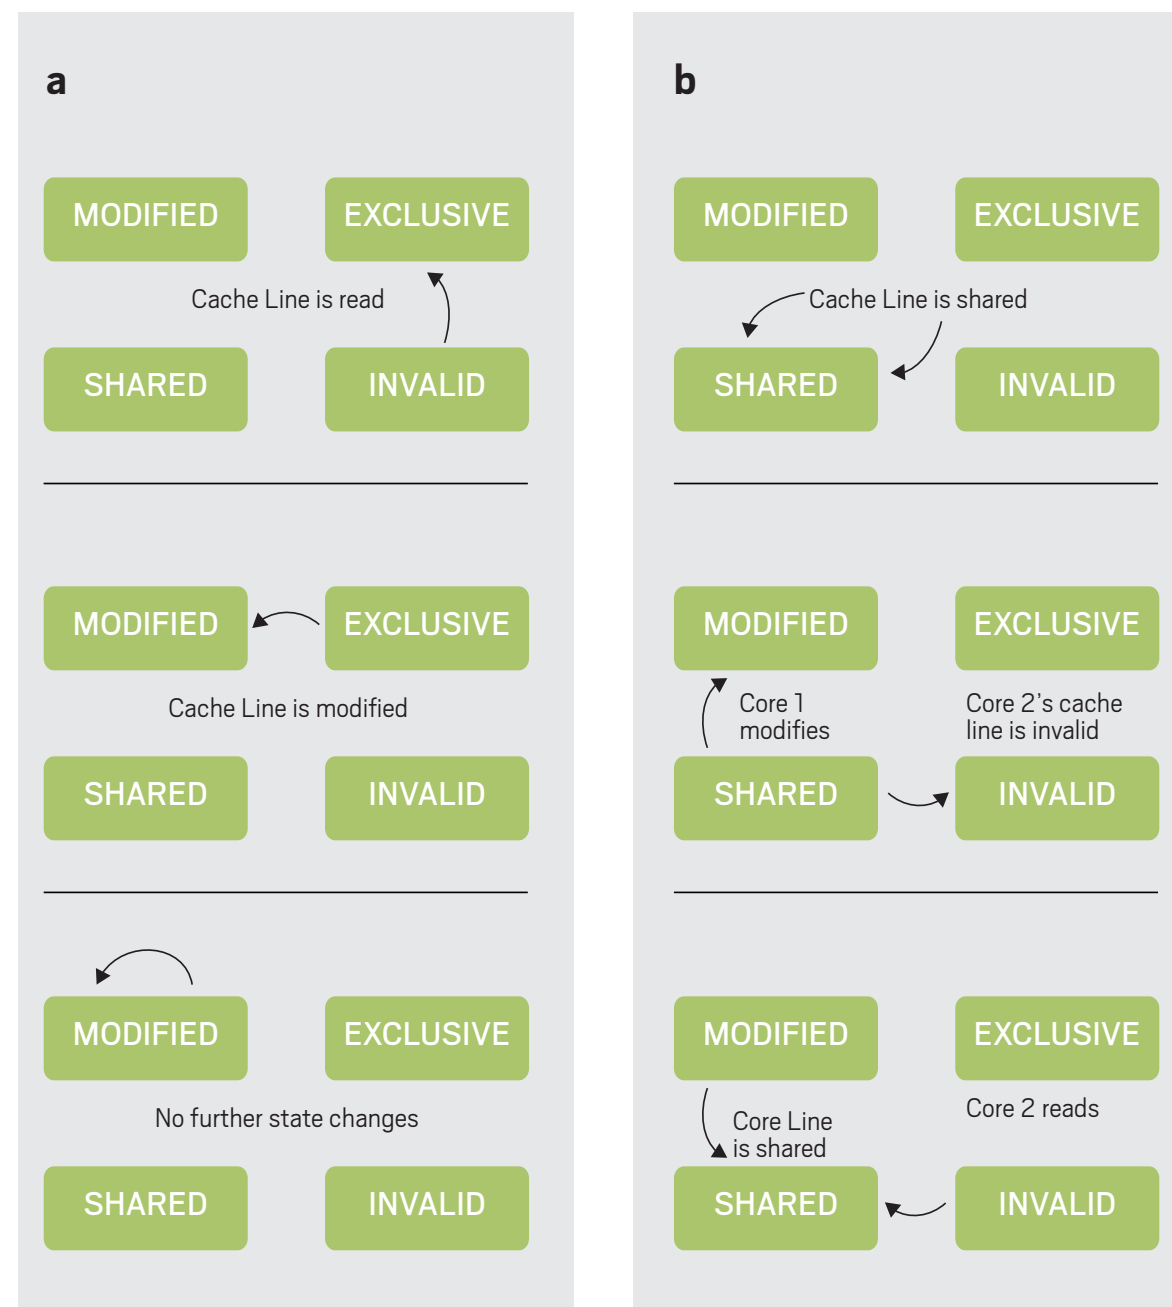
As each core copies the cache line into its respective L1 cache and does the update to the variable, the core notifies the other cores of the update and tells them to freshen the L1 cache copy they own.

Modern computers use a variant of the Modified Exclusive Shared Invalid (MESI) protocol to synchronize the L1 caches. The protocols used today have been improved, but the basic principle remains the same. In the MESI protocol, every cache line that was loaded into the cache is in one of four states: *Modified*, *Exclusive*, *Shared*, or *Invalid*. The Modified state means that the cache has an exclusive copy of the cache line and has modified it. If a cache line is in the Exclusive state, it also means that it is an exclusive copy, but it has not been modified yet. A cache line is in the Shared state if there are copies in two or more caches, but none of them has been modified. An Invalid cache line is no longer valid and cannot be used.

**Figure 1** shows the state changes during our experiment. **Figure 1a** shows the state changes during the

The MESI protocol ensures that **the system never works with two different versions of the same cache line.**

first test run, when the array elements were farther apart and the threads accessed different cache lines. When one of the threads loaded a cache line into the cache, the line was marked as being Exclusive. When the thread modified it, the state switched to Modified. From then on, the cache line remained in the Modified state. Access to the cache line was fast, because the Modified state tells us that we can update the cache line



**Figure 1: MESI state changes while modifying elements (1a: adjacent elements, 1b: distant elements)**





formance numbers are roughly the same, but there is a significant difference in the number of *branch misses*. [Screen shots with full perf results are available in the [download area](#). —Ed.] The linear search had 346.9 million branch misses, while the binary search had 526.4 million. What is a branch miss? To understand branch misses, we need to take a step back and take a look at a mechanism called *pipelining*.

**Pipelining.** When the CPU processes an instruction, it actually has several steps to perform. The instruction needs to be fetched from memory and decoded. That is, the CPU must figure out what kind of instruction it is dealing with. After that, the instruction needs to be executed and the result has to be written back to memory. **Figure 2** shows the internal structure of a CPU and how a single instruction A steps through the different stages on a mythical CPU with no pipeline.

The operations performed at each stage are quite different. The stages are implemented in different parts of the CPU. Thus, if a processor really worked as shown in **Figure 2**, it would be inefficient, because most of its parts would be idle most of the time, waiting for the next instruction to arrive. Early on, chip designers thought about ways to keep all components busy and came up with a solution called *pipelining*.

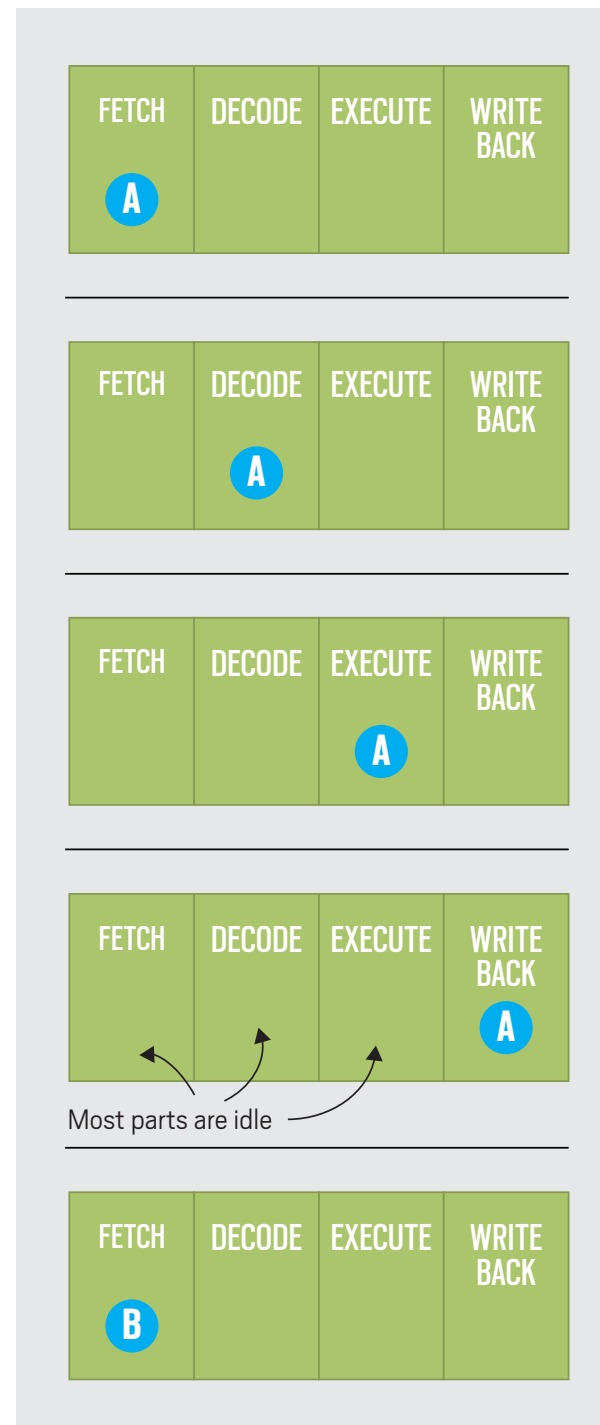
The principle can be seen in **Figure 3**. Instruction A is fetched. While instruction A is decoded, the CPU fetches the next instruction, B. As it executes instruction A, the CPU decodes instruction B, fetches the third instruction C, and so on. Instead of processing one instruction after another, the CPU processes a stream of instructions: the instruction pipeline.

This works extremely well. A single instruction still needs four cycles to step through all four stages, but with pipelining the throughput increases from one instruction every four cycles to one instruction per cycle.

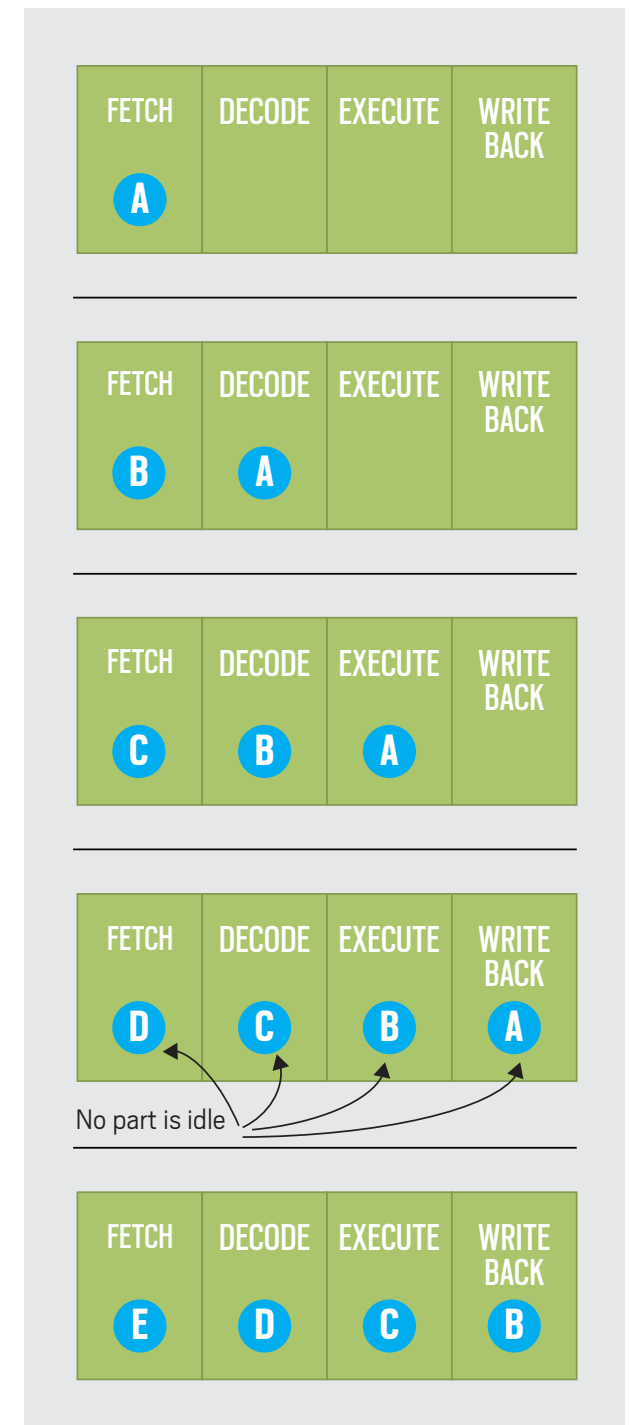
**Branch prediction.** There is one case, however, in which pipelining becomes tricky: conditional jumps. Imagine that instruction A is a conditional jump and that, depending on the outcome, it may execute instruction B or jump directly to a distant instruction X. Should the CPU fetch instruction B or X while we decode A? It does not know the outcome of the con-

ditional jump yet, because that becomes clear only after instruction A has finished the execution stage.

Different processor architectures deal with this situation differently. One of the most common strategies is *branch prediction*: The CPU guesses



**Figure 2:** CPU execution without pipeline



**Figure 3:** CPU execution with pipeline

In a typical program, the CPU's rate of **successful branch predictions** is well above **90 percent**.





# JSR 371: MVC 1.0 Specification

Specification Leads **Santiago Pericas-Geertsen** and **Manfred Riem** have joined forces with a collection of large and small organizations and individuals. If you'd like to comment on or contribute to JSR 371 or any other JSR, visit the [JCP website](#) and make your voice heard.

For all the details, check out [Ankara JUG's website](#), which is mostly in Turkish.





# JAVA CHAMPION PROFILE

## MARIO TORRE



**Mario Torre** has been an OpenJDK community member since its inception and is now an official member of the OpenJDK Adoption Group. In recent years, he has organized many of the famous Java DevRooms at FOSDEM.

**Java Magazine:** Where did you grow up?

**Torre:** In Salerno, a beautiful city in the south of Italy between Amalfi and the Cilento Coast. I now live in Hamburg, Germany.

**Java Magazine:** When and how did you first become

interested in computers?

**Torre:** When I was very young, my father bought me a Commodore 16. I was completely captured by it. I started to use PCs only when I was introduced to Linux, sometime in '97 or '98, although I owned a 386 for some time before that (which I used mostly to play X-Wing and Wing Commander).

**Java Magazine:** What was your first computer and programming language?

**Torre:** It was a BASIC dialect. I spent ages on the C16 getting the programs of the manual to work. I was introduced to C much later. One of my first real programs was a Norton Commander-like shell, which allowed users to organize files and programs. They could add shortcuts for up to nine of their favorite programs and launch them by simply pressing the associated number.

**Java Magazine:** What was your first professional programming job?

**Torre:** I always thought I would be a musician and programming was more of a hobby, so my first paid programming job came very late. It was Java-based and was something about organizing the lifecycle of an expedition, from arrival in the warehouse to delivery.

**Java Magazine:** What do you enjoy for fun and relaxation?

**Torre:** I usually play guitar and make music. But I also like model-making, drawing, reading, and creating my own music pedals. I'm working now on a digital delay on a Raspberry Pi.

**Java Magazine:** What “side effects” of your career do you enjoy the most?

**Torre:** Meeting people and traveling. I also deal a lot with students, introducing them to OpenJDK and the Java community. It's a lot of

work, but I really enjoy it.

**Java Magazine:** Has being a Java Champion changed anything for you with respect to your daily life?

**Torre:** I think that being a Java Champion is a way to recognize one's work. In that respect, you just keep doing the same things as before. But of course there is the personal side of it: I need to remember that what I say will have an even greater resonance than before.

**Java Magazine:** What are you looking forward to in the coming years?

**Torre:** I want to see what will happen next and contribute to it. I believe we are doing a great job of keeping OpenJDK vibrant. I intend to keep on pushing Oracle toward openness. It's great to be part of this change.

Follow [Mario Torre's blog](#)  
or reach him on Twitter  
[@neugens](#).





## Comments

We welcome your comments, corrections, opinions on topics we've covered, and any other thoughts you feel important to share with us or our readers. Unless you specifically tell us that your correspondence is private, we reserve the right to publish it in our Letters to the Editor section.

## Article Proposals

We welcome article proposals on all topics regarding Java and other JVM languages, as well as the JVM itself. We also are interested in proposals for articles on Java utilities (either open source

or those bundled with the JDK). Finally, algorithms, unusual but useful programming techniques, and most other topics that hard-core Java programmers would enjoy are of great interest to us, too. Please contact us with your ideas at [javamag\\_us@oracle.com](mailto:javamag_us@oracle.com) and we'll give you our thoughts on the topic and send you our nifty writer guidelines, which will give you more information on preparing an article.

## Customer Service

If you're having trouble with your subscription, please contact the folks at [java@halldata.com](mailto:java@halldata.com) (phone

+1.847.763.9635), who will do whatever they can to help.

# Where?

Comments and article proposals should be sent to me, **Andrew Binstock**, at [javamag\\_us@oracle.com](mailto:javamag_us@oracle.com).

While it will have no influence on our decision whether to publish your article or letter, cookies and edible treats will be gratefully accepted by our staff at *Java Magazine*, Oracle Corporation, 500 Oracle Parkway, MS OPL 3A, Redwood Shores, CA 94065, USA.

